

Deep importance sampling based on regression for model inversion and emulation



F. Llorente ^{a,*}, L. Martino ^b, D. Delgado-Gómez ^a, G. Camps-Valls ^c

^a Dep. of Statistics, Universidad Carlos III de Madrid, Leganés, Spain

^b Dep. of Signal Processing, Universidad Rey Juan Carlos, Móstoles, Spain

^c Dep. of Electrical Engineering, Universitat de València, Valencia, Spain

ARTICLE INFO

Article history:

Available online 20 May 2021

Keywords:

Model inversion
Bayesian inference
Emulation
Adaptive regression
Importance sampling
Remote sensing

ABSTRACT

Understanding systems by forward and inverse modeling is a recurrent topic of research in many domains of science and engineering. In this context, Monte Carlo methods have been widely used as powerful tools for numerical inference and optimization. They require the choice of a suitable proposal density that is crucial for their performance. For this reason, several adaptive importance sampling (AIS) schemes have been proposed in the literature. We here present an AIS framework called Regression-based Adaptive Deep Importance Sampling (RADIS). In RADIS, the key idea is the adaptive construction via regression of a non-parametric proposal density (i.e., an emulator), which mimics the posterior distribution and hence minimizes the mismatch between proposal and target densities. RADIS is based on a deep architecture of two (or more) nested IS schemes, in order to draw samples from the constructed emulator. The algorithm is highly efficient since employs the posterior approximation as proposal density, which can be improved adding more support points. As a consequence, RADIS asymptotically converges to an exact sampler under mild conditions. Additionally, the emulator produced by RADIS can be in turn used as a cheap surrogate model for further studies. We introduce two specific RADIS implementations that use Gaussian Processes (GPs) and Nearest Neighbors (NN) for constructing the emulator. Several numerical experiments and comparisons show the benefits of the proposed schemes. A real-world application in remote sensing model inversion and emulation confirms the validity of the approach.

© 2021 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Modeling and understanding systems is of paramount relevance in many domains of science and engineering. The problems involve both forward and inverse modeling, and very often one resorts to domain knowledge (either in the form of mechanistic models, hypotheses, constraints or just data) and observational data to learn parametrizations and do inferences. Among the many approaches possible, Bayesian methods have become very popular during the last decades. Bayesian inference is very active in the communities of machine learning, statistics and signal processing [1–3]. With them, there has been a surge of interest in the Monte Carlo (MC) techniques that are often necessary for the implementation of the Bayesian analysis. Several families of MC schemes have been proposed that excel in numerous applications, including the popular Markov Chain Monte Carlo (MCMC) algorithms, particle filtering techniques and adaptive importance sampling (AIS) methods [3,4].

Adaptive Importance Sampling (AIS) The performance of the MC algorithms depends strongly on the proper choice of a proposal probability density function (pdf). In adaptive schemes, the proposal pdf is updated considering the previous generated samples. In recent years, a plethora of AIS algorithms have been proposed in the literature [4]. In most of these algorithms, the complete proposal can be expressed as a finite parametric mixture of densities [5–9]. Unlike these schemes, we consider a non-parametric proposal based on an interpolating construction.

Emulators in Bayesian inference Furthermore, many Bayesian inference problems involve the evaluation of computationally intensive models, due to the use of particularly complex systems, consisting of many coupled ordinary or partial differential equations in high-dimensional spaces, or a large amount of available data. To overcome this issue, a successful approach consists in replacing the true model by a surrogate model (a.k.a. an emulator) [10–14].

The resulting emulator can be employed in different ways inside a Bayesian analysis. A first possibility is to apply MC sampling methods considering the surrogate model as an approximate posterior pdf within the MC schemes [15–17][18, Chapter 9.4.3] or

* Corresponding author.

E-mail address: felloren@est-econ.uc3m.es (F. Llorente).

within different quadrature rules [19–21], instead of the evaluation of a costly true posterior. For instance, this is also the case of the strategy known as *calibrate, emulate, sample*, currently in vogue [22]. In order to improve the efficiency of MC algorithms, a second option is to use the emulator as a proposal density within an MC technique. Here, we focus on the last approach.

Contribution In this work, we design a deep AIS framework where a non-parametric interpolating proposal density is adapted online. The new approach is called Regression-based Adaptive Deep Importance Sampling (RADIS). In RADIS, the key idea is the adaptive construction of a non-parametric proposal pdf (i.e., *an emulator*), which mimics the posterior distribution in order to minimize the mismatch between proposal and target pdfs. Differently from other adaptive schemes, the adaptation in RADIS not only uses the information of the previous samples, but also all the evaluations of the posterior for directly constructing the emulator. Thus, unlike in a parametric approach, in our setting this discrepancy can be arbitrarily decreased to zero by adding more nodes. Hence, RADIS is asymptotically an exact sampler. The proposed methodology is based on a *deep architecture*: two nested IS schemes are employed, with an inner and an outer IS layers. The inner IS stage is used to generate samples from the emulator. The outer IS layer provides the final posterior approximation by a cloud of weighted samples. Thus, RADIS finally provides two approximations of the posterior, one in form of a weighted particle measure, and also the emulator adapted online.¹ Parsimonious constructions of the emulator have been also discussed.

We discuss two specific implementation of RADIS. These specific implementations differ on the choice of the emulator construction. In the first one, a Gaussian Process (GP) model is applied to the log-posterior function obtaining the novel scheme denoted as GP-AIS. In the second one, a piece-wise constant approximation based on Nearest Neighbors (NNs) is applied, providing the novel algorithm denoted as NN-AIS. In both cases, the resulting proposal pdf can be seen as an *incremental mixture* of densities. A deep structure with more than two layers is described, where a chain of emulators is adapted and then employed as proposal pdfs within different nested IS stages. Robust and sequential implementations are also discussed. Several numerical comparisons show the advantages of RADIS with respect to benchmark algorithms. A real-word application illustrates the capabilities for sequential parameter retrieval and emulation of a well-known radiative transfer model (RTM) used in remote sensing. In the next section, a brief overview of the related approaches is provided.

2. Other related works

The non-parametric interpolating construction of the proposal and related strategies are appealing from different points of views. This is proved by attention devoted by the previous attempts in the literature shown above, and by other related approaches that we describe next.

Interpolating proposal The idea of using interpolating densities is particularly attractive since we can arbitrarily decrease the mismatch between proposal and posterior by adding more support points. For this reason, the resulting algorithms provide very good performance [23–27]. The first use of an interpolating procedure for building a proposal density can be ascribed to the rejection sampling and adaptive rejection sampling schemes [28–30]. The well-know Zigurat algorithm and table methods are other examples of fast rejection samplers employing interpolating proposals

[31,2]. They are state-of-the-art methods as random sample generators of specific univariate distributions in terms of speed of generation. In some rejection samplers and MCMC algorithms, the proposal is formed by polynomial pieces (constant, linear, etc.) [23–26], [2, Chapters 4 and 7]. The use of interpolating proposal pdfs within an IS scheme is also considered in [32]. The conditions needed for applying an emulator as a proposal density are discussed in [27]. More specifically, we need to be able to: **(a)** update the construction of the emulator, **(b)** evaluate the emulator, **(c)** normalize the function defined by the emulator, and **(d)** draw samples from the emulator. It is not straightforward to find an interpolating (or regression) construction which satisfies all those conditions jointly, and especially for an arbitrary dimensionality of the problem. This is the reason why the previous attempts of using an interpolating proposal pdfs are restricted to the univariate case. Our deep architecture solves these issues.

Partitioning and stratification Note also that the use of a proposal pdf formed by components restricted to disjoint regions of the domain (like in the piecewise constant proposal based on NN) is related to the stratification idea. Indeed, different schemes based on partitioning and/or stratification divide the entire domain in disjoint sub-regions and consider different partial proposals in each of them [3, Chapter 4.6.3], [33–36]. The complete proposal pdf is then a mixture of the partial proposals. Moreover, this process can be iterated so that the partition is refined over the iterations increasing the number of partial proposals. In this case, the complete proposal is an *incremental mixture* as RADIS (see also below) [33,36]. Recent works propose using trees in order to partition the space and subsequently build the proposal [32,37]. In the context of MCMC, [38] builds an approximation of the target using Polya trees.

Incremental mixtures The use of non-parametric but *non-interpolating* proposals have been suggested in other works. A non-parametric IS approach is considered in [39], where the proposal is built by a kernel density estimation. In [40], a proposal pdf defined as a mixture with increasing number of components is also suggested. When a weighting strategy based on the so-called *temporal deterministic mixture* is applied [41,42], incremental mixture proposals appear also in other IS schemes (e.g., [41,43]).

Other approaches Surrogate GP models has been also employed within IS schemes in the context of rare event estimation [44,45]. Finally, other IS schemes can be encompassed in a similar “deep” approach [46,47]. In the first one, MCMC steps are used to jump from different tempered versions of the posterior, and a global IS weighting as product of intermediate weights [46]. In the second scheme a two-stages weighting procedure is used, where the first layer considers a Gauss-Hermite quadrature and the second layer is a standard IS method [47].

3. Preliminaries and motivation

3.1. Problem statement

Bayesian inference In many real world applications, the goal is to infer a variable of interest given a set of data [1]. Let us denote the parameter of interest (static or dynamic) by $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^{d_x}$, and let $\mathbf{y} \in \mathbb{R}^{d_y}$ be the observed data. In a Bayesian analysis, all the statistical information is contained in the posterior distribution, which is given by

$$\bar{\pi}(\mathbf{x}) = p(\mathbf{x}|\mathbf{y}) = \frac{\ell(\mathbf{y}|\mathbf{x})g(\mathbf{x})}{Z(\mathbf{y})}, \quad (1)$$

¹ The emulation can be applied to the entire posterior or part of it, like a physical model.

where $\ell(\mathbf{y}|\mathbf{x})$ is the likelihood function, $g(\mathbf{x})$ is the prior pdf, and $Z(\mathbf{y})$ is the Bayesian model evidence (a.k.a. marginal likelihood). The marginal likelihood $Z(\mathbf{y})$ is important for model selection purposes [48,49]. Generally, $Z(\mathbf{y})$ is unknown, so we are able to evaluate the unnormalized target function, $\pi(\mathbf{x}) = \ell(\mathbf{y}|\mathbf{x})g(\mathbf{x})$. The analytical computation of the posterior density $\tilde{\pi}(\mathbf{x}) \propto \pi(\mathbf{x})$ is often unfeasible, hence numerical approximations are needed. Our goal is to approximate integrals of the form

$$I = \int_{\mathcal{X}} f(\mathbf{x})\tilde{\pi}(\mathbf{x})d\mathbf{x} = \frac{1}{Z} \int_{\mathcal{X}} f(\mathbf{x})\pi(\mathbf{x})d\mathbf{x}, \quad (2)$$

where $f(\mathbf{x})$ is some integrable function, and

$$Z = \int_{\mathcal{X}} \pi(\mathbf{x})d\mathbf{x}. \quad (3)$$

In the literature, random sampling or deterministic quadratures are often used [3,2,50]. In this work, we focus on the so-called IS approach.

Emulation There exist many situations where the evaluation of π is expensive (e.g., as in big data framework or when the observation model is costly). Hence, we are also interested in obtaining an emulator of $\pi(\mathbf{x})$ (or just a part of the posterior), denoted $\hat{\pi}_t(\mathbf{x})$, such that (i) $\hat{\pi}_t(\mathbf{x})$ is cheap to evaluate, and (ii) $\hat{\pi}_t(\mathbf{x}) \rightarrow \pi(\mathbf{x})$ (in some sense, e.g., L_2 norm) as $t \rightarrow \infty$.

3.2. Importance sampling (IS) and aim of the work

Let us consider a normalized proposal density $\bar{q}(\mathbf{x})$.² The importance sampling (IS) method consists of drawing N independent samples, $\mathbf{x}_1, \dots, \mathbf{x}_N$, from $\bar{q}(\mathbf{x})$ (also called particles), and then assign to each sample the following unnormalized weights

$$w_n = w(\mathbf{x}_n) = \frac{\pi(\mathbf{x}_n)}{\bar{q}(\mathbf{x}_n)}, \quad n = 1, \dots, N. \quad (4)$$

An unbiased estimator of the marginal likelihood Z is given by the arithmetic mean of these unnormalized weights [51,3], i.e.,

$$\hat{Z} = \frac{1}{N} \sum_{n=1}^N w_n.$$

Defining also the normalized weights $\bar{w}_n = \frac{w_n}{\sum_{i=1}^N w_i}$, with $n = 1, \dots, N$, the self-normalized IS estimator of I in Eq. (2) is given by

$$\hat{I} = \sum_{n=1}^N \bar{w}_n f(\mathbf{x}_n).$$

More generally, regardless of the specific function $f(\mathbf{x})$, we obtain a particle approximation of $\tilde{\pi}$, i.e., $\tilde{\pi}(\mathbf{x}) = \sum_{n=1}^N \bar{w}_n \delta(\mathbf{x} - \mathbf{x}_n)$, where $\delta(\mathbf{x})$ is a delta function. It is important to remark that with this particle approximation, we can approximate several quantities related to the posterior $\tilde{\pi}(\mathbf{x})$, such as any moments and/or credible intervals (not just a specific integral). The quality of this particle approximation is related to the discrepancy between the proposal $\bar{q}(\mathbf{x})$ and the posterior $\tilde{\pi}(\mathbf{x})$. Indeed, in an ideal MC scenario, we can draw from the posterior, i.e., $\bar{q}(\mathbf{x}) = \tilde{\pi}(\mathbf{x})$, so that $\bar{w}_n = \frac{1}{N}$, which corresponds with the maximum effective sample

size (ESS) [51,52]. With a generic proposal $\bar{q}(\mathbf{x}) \propto \tilde{\pi}(\mathbf{x})$, we can obtain a very small ESS and a bad particle approximation $\tilde{\pi}(\mathbf{x})$ (i.e., poor performance of the algorithm).

Remark 1. The variance of the marginal likelihood estimator $\hat{Z} = \frac{1}{N} \sum_{n=1}^N w(\mathbf{x}_n)$ is given by

$$\text{var}[\hat{Z}] = \frac{1}{N} \text{var}[w(\mathbf{x})], \quad (5)$$

where $w(\mathbf{x}) = \frac{\pi(\mathbf{x})}{\bar{q}(\mathbf{x})}$ and $\mathbf{x} \sim \bar{q}(\mathbf{x})$. Since \hat{Z} is also unbiased, then we also have

$$\mathbb{E}[|Z - \hat{Z}|^2] = \frac{1}{N} \text{var}[w(\mathbf{x})]. \quad (6)$$

For more details, see [3].

Remark 2. The variance of the IS weight function $w(\mathbf{x})$ is proportional to the Pearson divergence between $\bar{q}(\mathbf{x})$ and $\tilde{\pi}(\mathbf{x})$, denoted as $\chi^2(\tilde{\pi} \parallel \bar{q})$ (also called χ^2 distance), i.e.,

$$\text{var}[w(\mathbf{x})] \propto \chi^2(\tilde{\pi} \parallel \bar{q}) = \int_{\mathcal{X}} \frac{(\tilde{\pi}(\mathbf{x}) - \bar{q}(\mathbf{x}))^2}{\bar{q}(\mathbf{x})} d\mathbf{x}. \quad (7)$$

See [52,53] and Appendix A.2 for further details. Regarding the mean squared error of the estimator \hat{I} , we have

$$\mathbb{E}[|I - \hat{I}|^2] \leq \frac{C_f}{N} (\chi^2(\tilde{\pi} \parallel \bar{q}) + 1). \quad (8)$$

The relationships with the L_2 and L_∞ distances is also given in Appendix A.2.

To reduce the discrepancy between the proposal $\bar{q}(\mathbf{x})$ and the posterior $\tilde{\pi}(\mathbf{x})$, we consider a non-parametric adaptive construction of the proposal $\bar{q}_t(\mathbf{x})$ where t denotes a discrete iteration index. In order to make that the discrepancy becomes smaller and smaller, an interpolating procedure $\bar{q}_t(\mathbf{x})$ based on a set of support points \mathcal{S}_{t-1} is employed. Namely, we generate a sequence of proposal pdfs $\bar{q}_t(\mathbf{x}), \bar{q}_{t+1}(\mathbf{x}), \bar{q}_{t+2}(\mathbf{x}), \dots$ which become closer and closer to $\tilde{\pi}(\mathbf{x})$, as the number of support points grows. Throughout the paper, we denote $\bar{q}_t(\mathbf{x}) \propto \hat{\pi}_t(\mathbf{x})$ the non-parametric regression function which approximates the unnormalized posterior $\pi(\mathbf{x})$ at iteration t . The normalized proposal is denoted as $\bar{q}_t(\mathbf{x}) = \frac{1}{c_t} \hat{\pi}_t(\mathbf{x})$, where $c_t = \int_{\mathcal{X}} \hat{\pi}_t(\mathbf{x}) d\mathbf{x}$. Although the approximation $\hat{\pi}_t$ depends on the set of nodes \mathcal{S}_{t-1} , for simplicity we use the simpler notation $\hat{\pi}_t(\mathbf{x}) = \hat{\pi}_t(\mathbf{x}; \mathcal{S}_{t-1})$.

Remark 3. If the sequence of proposals is such $\|\tilde{\pi} - \bar{q}_t\|_2 \rightarrow 0$ as $t \rightarrow \infty$, then $\chi^2(\tilde{\pi} \parallel \bar{q}_t) \rightarrow 0$. See Appendix A.2 for more details.

4. Regression-based Adaptive Deep Importance Sampling

In this section, we introduce the proposed scheme, called Regression-based Adaptive Deep Importance Sampling (RADIS). The resulting algorithm is an adaptive importance sampler with a non-parametric interpolating proposal pdf. We show how to implement the sampling and construction of the proposal density in Sect. 4.1 and Sect. 4.2 respectively. The novel scheme is summarized in Table 1. The proposal is adaptively built using a regression approach that considers the set of all previous nodes \mathbf{x}_i 's where π is evaluated. In Section 4.2, we present two construction methodologies considered in this work. Samples from this proposal are drawn via an approximate procedure that can be interpreted as an additional "inner" IS. All the samples generated in the inner IS are then used in the "outer" IS. Fig. 1 outlines this procedure. The adaptation consists in sequentially adding the samples to the set of current nodes

² We assume that $\bar{q}(\mathbf{x}) > 0$ for all \mathbf{x} where $\tilde{\pi}(\mathbf{x}) > 0$, and $\bar{q}(\mathbf{x})$ has heavier tails than $\tilde{\pi}(\mathbf{x})$.

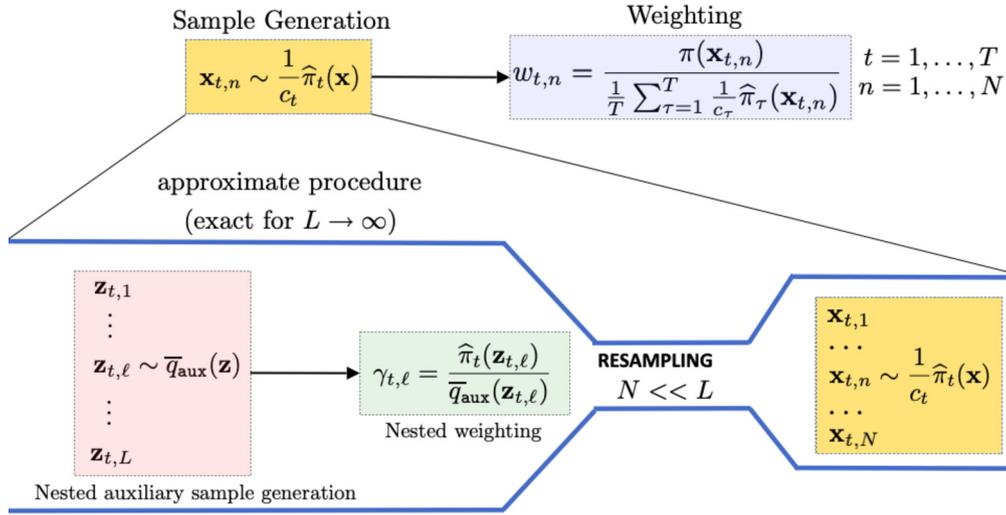


Fig. 1. Approximate sampling from $\frac{1}{c_t} \hat{\pi}_t(\mathbf{x}) \propto \hat{\pi}_t(\mathbf{x})$ and final weighting scheme.

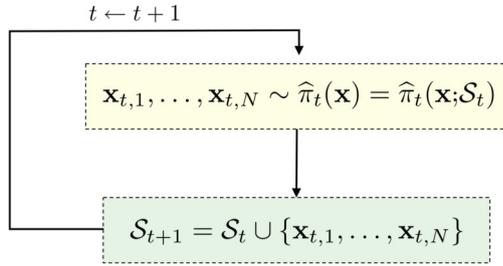


Fig. 2. Graphical representation of the adaptation scheme. More parsimonious alternatives are introduced in Section 6.

(see Fig. 2). In the outer IS, we consider a temporal deterministic mixture approach to compute the weights. Note that the weighting step needs to be done only once at the end of the algorithm.

4.1. RADIS: a two-layer Deep IS

RADIS is an adaptive IS scheme based on two IS stages. In the following, we describe the inner and outer stages as well as the possible construction and adaptation of the non-parametric proposal density. The extension with more than two nested layers is also discussed.

4.1.1. Inner IS scheme

The inner IS stage is repeated at every iteration. It generates samples approximately distributed from the current non-parametric proposal, denoted as $\hat{\pi}_t$ (the unnormalized version). Furthermore, these samples are used to normalize $\hat{\pi}_t$, i.e., in order to estimate $c_t = \int_{\mathcal{X}} \hat{\pi}_t(\mathbf{x}) d\mathbf{x}$.

Approximate sampling from the emulator It is not straightforward to sample from an interpolating proposal [23,27]. We propose using an approximate procedure based on IS. Specifically, at each iteration, in order to sample from $\frac{1}{c_t} \hat{\pi}_t(\mathbf{x})$, we use sampling importance resampling (SIR) with an auxiliary proposal \bar{q}_{aux} [54]. First, a set of $\{\mathbf{z}_{t,\ell}\}_{\ell=1}^L$ (with large L) are drawn from $\bar{q}_{\text{aux}}(\mathbf{x})$. These auxiliary samples are weighted according to $\hat{\pi}_t(\mathbf{x})$

$$\gamma_{t,\ell} = \frac{\hat{\pi}_t(\mathbf{z}_{t,\ell})}{\bar{q}_{\text{aux}}(\mathbf{z}_{t,\ell})} \quad \ell = 1, \dots, L.$$

Finally, in order to obtain $\{\mathbf{x}_{t,n}\}_{n=1}^N$, we resample N times within $\{\mathbf{z}_{t,\ell}\}_{\ell=1}^L$ with probabilities $\{\tilde{\gamma}_{t,\ell}\}_{\ell=1}^L$ where $\tilde{\gamma}_{t,\ell} = \frac{\gamma_{t,\ell}}{\sum_{i=1}^L \gamma_{t,i}}$ for $\ell = 1, \dots, L$, i.e.,

$$\mathbf{x}_{t,n} \sim \sum_{\ell=1}^L \tilde{\gamma}_{t,\ell} \delta(\mathbf{x} - \mathbf{z}_{t,\ell}), \quad \text{for all } n. \quad (9)$$

In this way, we obtain a set of samples $\{\mathbf{x}_{t,n}\}_{n=1}^N$ approximately distributed from $\hat{\pi}_t$ [54,55].

Remark 4. Under some mild conditions, as $L \rightarrow \infty$, the SIR procedure is asymptotically exact. Namely, as $L \rightarrow \infty$ the density of the resampled particles becomes closer and closer to $q_t(\mathbf{x}) \propto \hat{\pi}_t(\mathbf{x})$. See, for instance, the following references [54], [56, Sect. 6.2.4], [55, Sect. 3.2]. For further details, see [57, page 6], [58, App. A] and also Appendix A.

Remark 5. Note that the computation of the inner IS weights $\gamma_{t,\ell}$'s does not involve the evaluation of the posterior $\pi(\mathbf{x})$, but only the evaluation of the emulator $\hat{\pi}_t(\mathbf{x})$. Hence, assuming that the evaluation of the posterior is the main computational bottleneck, in this setting we can make L arbitrarily large.

Since we resample from a finite set, we can obtain duplicated samples, but it rarely happens when $L \gg N$. An alternative to avoid these repetitions is to use a regularized resampling, i.e.,

$$\mathbf{x}_{t,n} \sim \sum_{\ell=1}^L \tilde{\gamma}_{t,\ell} K(\mathbf{x} - \mathbf{z}_{t,\ell}), \quad \text{for all } n, \quad (10)$$

where the Dirac measures have been replaced by a kernel function $K(\mathbf{x})$ [59]. The bandwidth of $K(\mathbf{x})$ can be tuned according to some kernel density estimation (KDE) criterion. For the computation of the outer IS weights (see below), we need to approximate $c_t = \int_{\mathcal{X}} \hat{\pi}_t(\mathbf{x}) d\mathbf{x}$ for $t = 1, \dots, T$. They are estimated during the inner IS by the corresponding estimator, $\hat{c}_t = \frac{1}{L} \sum_{\ell=1}^L \gamma_{t,\ell}$, for $t = 1, \dots, T$. We have $\hat{c}_t \rightarrow c_t$ when $L \rightarrow \infty$, by standard IS arguments [3].

4.1.2. Adaptation

At each iteration, at the end of the inner IS stage, the algorithm performs the adaptation producing $\hat{\pi}_{t+1}$. Specifically, the emulator $\hat{\pi}_t(\mathbf{x})$ is improved by incorporating the generated samples at each iteration as additional nodes (see Fig. 2). Namely, the additional support points $\{\mathbf{x}_{t,n}\}_{n=1}^N$ to \mathcal{S}_t are obtained by resampling N

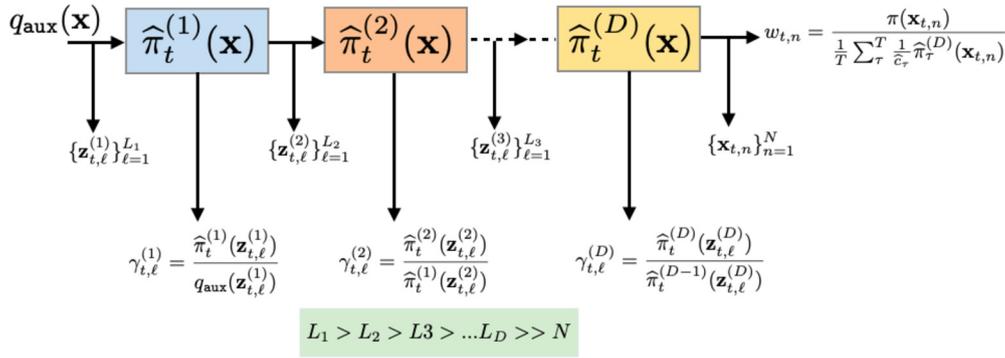


Fig. 3. RADIS with $D + 1$ layers in the deep architecture. Different emulator construction can be applied at each stage. In each d -th layer, the resampling is applied L_{d+1} times for generating the next cloud of resampled particles $\{\mathbf{z}_{t,\ell}^{(d+1)}\}_{\ell=1}^{L_{d+1}}$ (with $d = 1, \dots, D + 1$). These samples are used for the adaptation of $\hat{\pi}_t^{(d)}$ and then are weighted again in the next stage. Note that $L_{D+1} = N$ and $L_d > L_{d+1}$.

times within $\{\mathbf{z}_{t,\ell}\}_{\ell=1}^L$ according to the probabilities $\tilde{\gamma}_{t,\ell} = \frac{\gamma_{t,\ell}}{\sum_{i=1}^L \gamma_{t,i}}$ for $\ell = 1, \dots, L$. Note that the probability mass $\tilde{\gamma}_{t,\ell}$ is directly proportional to $\hat{\pi}_t(\mathbf{z}_{t,\ell})$. Therefore, the algorithm tends to add points where $\hat{\pi}_t$ is higher. Indeed, as $L \rightarrow \infty$, the resampled particles are distributed as $\hat{\pi}_t$ [54–56]. If L is not great enough, some $\mathbf{x}_{t,n}$ can be repeated. We do not include these repetitions as support points. Increasing L or using a regularized resampling as in Eq. (10) avoids this issue [59]. Note that the number of support points $J_t = |\mathcal{S}_t|$ increases as t grows.

All the evaluations of the unnormalized posterior $\pi(\mathbf{x})$ in the additional nodes are stored in the vector denoted as $\boldsymbol{\pi}_t$, in order to be used in the outer IS stage. Note also that all evaluations of π are used to build the emulator.

4.1.3. Outer IS scheme

At the end of the iterative part, we compute the final IS weights $w_{t,n}$, using all the posterior evaluations $\pi_{t,n} = \pi(\mathbf{x}_{t,n})$, which are stored in the inner layer. More specifically, we assign to each sample (drawn also in the inner stage) the weight

$$w_{t,n} = \frac{\pi_{t,n}}{\frac{1}{T} \sum_{\tau=1}^T \frac{1}{c_\tau} \hat{\pi}_\tau(\mathbf{x}_{t,n})}, \quad \text{for all } t = 1, \dots, T, \quad n = 1, \dots, N, \quad (11)$$

where we have employed a deterministic mixture weighting scheme [60,42], i.e., the denominator consists of a temporal mixture (e.g., as also suggested in [43]). Note that the weights $w_{t,n}$ are not required in the iterative inner layer described above. Hence, they can be computed after the adaptation and sampling steps are finalized. The output of the algorithm is then formed by all the sets of weighted particles $\{\mathbf{x}_{t,n}, w_{t,n}\}_{n=1}^N$ for $t = 1, \dots, T$, and the final emulator $\hat{\pi}_{T+1}(\mathbf{x}) = \hat{\pi}_{T+1}(\mathbf{x}; \mathcal{S}_T)$.

Remark 6. As $t \rightarrow \infty$ and $L \rightarrow \infty$, then $\hat{c}_t \rightarrow c_t \rightarrow Z$, i.e., is an approximation of the marginal likelihood. Another estimator of the marginal likelihood Z provided by RADIS is the arithmetic mean of all the outer weights, i.e., $\bar{Z} = \frac{1}{NT} \sum_{t=1}^T \sum_{n=1}^N w_{t,n}$.

Remark 7. Additional layers can be included in the proposed deep architecture would consists in adapting a chain of several emulators. This is graphically represented in Fig. 3. One of the advantages of this deep approach with $D + 1 > 2$ layers (where D is the number of inner nested stages), is that different emulator constructions can be jointly applied. Each emulator serves as proposal of the next IS stage. In the additional layers, the evaluation of the posterior (true model) is not required. In this scenario, RADIS also provides D different emulators.

Table 1
Regression-based Adaptive Deep Importance Sampling (RADIS).

<p>- Initialization: Choose the initial set \mathcal{S}_0 of nodes, and the values T, L, N (with $L \gg N$). Obtain the vector of initial evaluations $\boldsymbol{\pi}_0$.</p> <p>- For $t = 1, \dots, T$:</p> <ol style="list-style-type: none"> 1. Emulator construction: Given the set \mathcal{S}_{t-1} and the corresponding vector of posterior evaluations $\boldsymbol{\pi}_{t-1}$, build the proposal function $\hat{\pi}_t(\mathbf{x}) = \hat{\pi}_t(\mathbf{x} \mathcal{S}_{t-1})$ with a non-parametric regression procedure (see Sect. 4.2). 2. Inner IS: <ol style="list-style-type: none"> (a) IS. Sample $\{\mathbf{z}_{t,\ell}\}_{\ell=1}^L \sim q_{\text{aux}}(\mathbf{x})$ and compute the following weights $\gamma_{t,\ell} = \frac{\hat{\pi}_t(\mathbf{z}_{t,\ell})}{q_{\text{aux}}(\mathbf{z}_{t,\ell})}, \quad (12)$ for $\ell = 1, \dots, L$. (b) Resampling. Resample $\{\mathbf{x}_{t,n}\}_{n=1}^N$ from $\{\mathbf{z}_{t,\ell}\}_{\ell=1}^L$ with probabilities $\{\tilde{\gamma}_{t,\ell}\}_{\ell=1}^L$ where $\tilde{\gamma}_{t,\ell} = \frac{\gamma_{t,\ell}}{\sum_{i=1}^L \gamma_{t,i}}$ for $\ell = 1, \dots, L$. (c) Normalizing constant. Compute $\hat{c}_t = \frac{1}{L} \sum_{\ell=1}^L \gamma_{t,\ell}. \quad (13)$ 3. Update: Evaluate $\pi_{t,n} = \pi(\mathbf{x}_{t,n})$, for all $n = 1, \dots, N$, and update the set of nodes appending $\mathcal{S}_t = \mathcal{S}_{t-1} \cup \{\mathbf{x}_{t,1}, \dots, \mathbf{x}_{t,N}\}$ and $\boldsymbol{\pi}_t = [\boldsymbol{\pi}_{t-1}, \pi_{t,1}, \dots, \pi_{t,N}]^\top$. <p>- Outer IS: Assign to each sample the weight</p> $w_{t,n} = \frac{\pi_{t,n}}{\frac{1}{T} \sum_{\tau=1}^T \frac{1}{c_\tau} \hat{\pi}_\tau(\mathbf{x}_{t,n})}, \quad \text{for all } t = 1, \dots, T, \quad n = 1, \dots, N.$ <p>- Outputs: Final emulator $\hat{\pi}_{T+1}(\mathbf{x}) = \hat{\pi}_{T+1}(\mathbf{x} \mathcal{S}_T)$, and the set of weighted particles $\{\mathbf{x}_{t,n}, w_{t,n}\}_{n=1}^N$ for $t = 1, \dots, T$.</p>
--

4.2. Construction of $\hat{\pi}$ by regression

We consider two different procedures to build the non-parametric proposal: a Gaussian process (GP) model and nearest neighbors (NN) scheme. In Appendix A.4, we show that these constructions converges to the true underlying function as the number of nodes ($J_t = |\mathcal{S}_t|$) grows.

GP construction Let us consider building the surrogate $\hat{\pi}$ with Gaussian process (GP) regression in the log domain, i.e., over the $\log \pi(\mathbf{x})$ [61,23]. GP regression provides with an approximation of a function from a set $\mathbf{x}_1, \dots, \mathbf{x}_{J_t} \in \mathcal{X} \subseteq \mathbb{R}^{d_x}$ (where \mathcal{X} can be unbounded) and their corresponding function evaluation [61,62]. To ensure the non-negativity of the approximation, we fit the GP to $\log \pi$ rather than directly on π [63]. Let $\phi(\mathbf{x}) \equiv \log \pi(\mathbf{x})$ and $\boldsymbol{\phi} = [\phi_1, \dots, \phi_{J_t}]^\top$ where $\phi_i = \log \pi(\mathbf{x}_i)$ for $i = 1, \dots, J_t$. Given a symmetric and positive definite kernel $k(\mathbf{x}, \mathbf{x}')$ and some noise

level σ , under the assumption that $\phi(\mathbf{x})$ is a zero-mean GP with kernel k , the GP regression of $\phi(\mathbf{x})$ is of the form

$$\widehat{\phi}_t(\mathbf{x}) = \sum_{i=1}^{J_t} \beta_i k(\mathbf{x}, \mathbf{x}_i), \quad (14)$$

where the coefficients $\boldsymbol{\beta} = [\beta_1, \dots, \beta_{J_t}]^\top$ are given by

$$\boldsymbol{\beta} = (\mathbf{K} + \zeta \mathbf{I})^{-1} \boldsymbol{\phi} \quad (15)$$

with $(\mathbf{K})_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$ for $1 \leq i, j \leq J_t$ and \mathbf{I} is the $J_t \times J_t$ identity matrix. Note that, for $\zeta = 0$, $\widehat{\phi}$ corresponds to an interpolator of ϕ . Note also that the cost of obtaining $\widehat{\phi}$ is $\mathcal{O}(J_t^3)$ since it requires inverting a $J_t \times J_t$ matrix. As an example, a possible choice of kernel is the Gaussian $k(\mathbf{x}, \mathbf{x}') = \exp\{-\frac{1}{2\epsilon^2} \|\mathbf{x} - \mathbf{x}'\|_2^2\}$, where the hyperparameter ϵ can be estimated, e.g., by maximizing the marginal likelihood [61]. Finally, the approximation of π is given by

$$\widehat{\pi}_t(\mathbf{x}) = \exp\{\widehat{\phi}_t(\mathbf{x})\}. \quad (16)$$

Instead of building on the emulator in the log-domain, a simpler alternative (to ensure non-negativity) consists in setting $\phi(\mathbf{x}) \equiv \pi(\mathbf{x})$, $\boldsymbol{\phi} = [\phi_1, \dots, \phi_{J_t}]^\top$ where $\phi_i = \pi(\mathbf{x}_i)$ for $i = 1, \dots, J_t$. Then, we set again $\boldsymbol{\beta} = [\beta_1, \dots, \beta_{J_t}] = (\mathbf{K} + \zeta \mathbf{I})^{-1} \boldsymbol{\phi}$ and $\widehat{\phi}_t(\mathbf{x}) = \sum_{i=1}^{J_t} \beta_i k(\mathbf{x}, \mathbf{x}_i)$. The emulator is finally obtained as

$$\widehat{\pi}_t(\mathbf{x}) = \max\{\widehat{\phi}_t(\mathbf{x}), 0\}. \quad (17)$$

Note that these approximations can be directly applied for unbounded support \mathcal{X} . We call the scheme based on these constructions as Gaussian Process Adaptive Importance Sampling (GP-AIS).

NN construction Given $\mathbf{x}_1, \dots, \mathbf{x}_{J_t} \in \mathcal{X} \subset \mathbb{R}^{d_x}$ (where \mathcal{X} is bounded) and evaluations $\pi(\mathbf{x}_1), \dots, \pi(\mathbf{x}_{J_t})$, the nearest neighbor (NN) interpolator at \mathbf{x} consists of assigning the value of its nearest node. This is equivalent to consider the Voronoi partition $\mathcal{X} = \cup_{i=1}^{J_t} \mathcal{R}_i$, where

$$\mathcal{R}_i = \{\mathbf{x} \in \mathcal{X} : \|\mathbf{x} - \mathbf{x}_i\| < \|\mathbf{x} - \mathbf{x}_j\| \text{ for } j \neq i\}, \quad (18)$$

is the i -th Voronoi cell. The NN interpolator of π is then given by

$$\widehat{\pi}_t(\mathbf{x}) = \sum_{i=1}^{J_t} \pi(\mathbf{x}_i) \mathbb{I}_{\mathcal{R}_i}(\mathbf{x}), \quad \mathbf{x} \in \mathcal{X}, \quad (19)$$

where $\mathbb{I}_{\mathcal{R}_i}(\mathbf{x})$ is the indicator function in \mathcal{R}_i . Note that $\widehat{\pi}$ above is an interpolating approximation of π . The NN search has a cost of $\mathcal{O}(J_t)$. We denote the scheme based on this construction as Nearest Neighbor Adaptive Importance Sampling (NN-AIS). The regression case consists in considering the k nearest neighbors to \mathbf{x} , and taking the arithmetic mean of the values π in those k nearest nodes.

Remark 8. Note that RADIS employs an incremental mixture proposal density. Indeed, the emulator $\widehat{\pi}_t(\mathbf{x})$ in Eqs. (14)-(16) and (19) can be expressed as a mixture of pdfs where the number of components, J_t , increases as t grows. For more details of the NN case, see Appendix B.

Remark 9. Under mild conditions, the emulator $\widehat{\pi}_t \rightarrow \pi$ and $\widehat{c}_t \rightarrow Z$ as $t \rightarrow \infty$ (and $L \rightarrow \infty$), hence $\frac{1}{c_t} \widehat{\pi}_t \rightarrow \bar{\pi}$ (see Appendix A.4). Moreover, the SIR scheme to draw from $\frac{1}{c_t} \widehat{\pi}_t$ is asymptotically exact when $L \rightarrow \infty$ (see Appendix A.1). Hence, RADIS is drawing samples from $\bar{\pi}$, i.e., it is asymptotically an exact sampler.

The GP construction provides smoother solutions that can be directly employed in unbounded domains. However, the GP requires the inversion of matrix (with a dimension that increases as the number of nodes grows) and the tuning of the hyperparameters of the kernel function. In contrast, the NN construction does not need any matrix inversion and, if we fix in advance the number k neighbors (for instance in the interpolation case, we have $k = 1$) no hyperparameter tuning is required.

5. Robust accelerating schemes

In this section, we present some alternatives in order to **(a)** reduce the dependence from the initial nodes and **(b)** increase the applicability of RADIS, **(c)** speed up the convergence of the emulator covering quickly the state space and finally **(d)** we discuss the computational cost of the proposed overall scheme. The resulting methods are robust schemes, which can be also employed for extending the use of NN-AIS in unbounded supports. This is achieved combining the non-parametric proposal function $\widehat{\pi}_t(\mathbf{x})$ with a parametric proposal density, $q_{\text{par}}(\mathbf{x})$. Hence, the complete proposal, denoted as $\varphi_t(\mathbf{x})$, will be a mixture of densities with a parametric and a non-parametric components.

Mixture with parametric proposal The use of an additional parametric density $\bar{q}_{\text{par}}(\mathbf{x})$ can (i) ensure that the complete proposal has fatter tails than target pdf, and (ii) foster the exploration of important regions that could be initially ignored due to a possible bad initialization. Thus, we consider the following mixture as a proposal density in the inner IS layer,

$$\varphi_t(\mathbf{x}) = \alpha_t \bar{q}_{\text{par}}(\mathbf{x}) + (1 - \alpha_t) \frac{1}{c_t} \widehat{\pi}_t(\mathbf{x}), \quad (20)$$

where $\alpha_t \in [0, 1]$ for all t , and α_t is a non-increasing function t . The idea is to set initially $\alpha_0 = \frac{1}{2}$, and then decrease $\alpha_t \rightarrow \alpha_\infty$ as $t \rightarrow \infty$ (e.g., we can set $\alpha_\infty = 0$). Note that $\varphi_t(\mathbf{x})$ must be evaluated in the denominator of the outer layer weights $w_{t,n}$ in (11), taking the place of $\frac{1}{c_t} \widehat{\pi}_t(\mathbf{x})$ (see Table 1).

Remark 10. Choosing $\bar{q}_{\text{par}}(\mathbf{x})$ with fatter tails than $\bar{\pi}(x)$, then $\varphi_t(\mathbf{x})$ has also fatter tails than $\bar{\pi}(x)$. Hence, we avoid the infinite variance issue of the IS weights [3].

See also [48, Section 7.1] for a theoretical and numerical example of the infinite variance problem. As an example, if \mathcal{X} is bounded, $\bar{q}_{\text{par}}(\mathbf{x})$ could be a uniform density over \mathcal{X} . If \mathcal{X} is unbounded, $\bar{q}_{\text{par}}(\mathbf{x})$ can be, e.g., a Gaussian, a Student-t distribution or a mixture of pdfs (see below).

Remark 11. The fact that $\varphi_t(\mathbf{x})$ has fatter tails than $\bar{\pi}(x)$ ensures to have a non-zero probability of adding new nodes in any possible subset of the support \mathcal{X} .

This strategy also allows the use of the NN-AIS in an unbounded support. In Appendix C we describe an extension of NN-AIS where the support of the NN approximation is also adapted.

Parametric mixture by other AIS schemes A more sophisticated option is to also update $\bar{q}_{\text{par}}(\mathbf{x})$ along the iterations. For instance, $\bar{q}_{\text{par}}(\mathbf{x}) = \frac{1}{C} \sum_{c=1}^C q_c(\mathbf{x} | \boldsymbol{\mu}_{t,c}, \boldsymbol{\Sigma}_{t,c})$ can be itself a mixture, whose parameters are adapted following another AIS scheme, so that the complete proposal would be

$$\varphi_t(\mathbf{x}) = \alpha_t \left(\frac{1}{C} \sum_{c=1}^C \bar{q}_c(\mathbf{x} | \boldsymbol{\mu}_{t,c}, \boldsymbol{\Sigma}_{t,c}) \right) + (1 - \alpha_t) \frac{1}{c_t} \widehat{\pi}_t(\mathbf{x}), \quad (21)$$

with $\alpha_t \in [0, 1]$ for all t . As an example, the parametric mixture $\bar{q}_{\text{par}}(\mathbf{x})$ can be obtained following a population Monte Carlo (PMC) method, or a layered adaptive importance sampling (LAIS) technique and/or adaptive multiple importance sampling (AMIS) scheme [4]. The weight α_t is again a non-increasing function of the iteration t .

Regression versus interpolation In the first iterations of RADIS, the use of $\zeta > 0$ in the GP approximation and/or considering the k nearest neighbors (instead only the closest one, $k = 1$), also decreases the dependence on the initial nodes. Namely, reducing the overfitting, at least in the first iteration of RADIS, also increases the robustness of the algorithm.

More layers To leverage the benefits of different emulator constructions in RADIS, one possible strategy is to employ additional layers in the deep architecture, as depicted in Fig. 3. For instance, with one additional layer, we could use jointly the GP and the NN constructions. Another possibility is to consider several GP models with different kernel functions or several NN schemes with different k .

5.1. Computational cost

In this section, we discuss computational details of our approach and hypothesize when our approach is convenient also in terms of computational time. It is important to remark that RADIS is useful also for constructing a good emulator (not just for approximating integrals as other Monte Carlo schemes), choosing the nodes in a proper way, similarly in an active learning scheme [21,14]. Figs. 6(d) and 11 in the numerical experiments provide a comparison with a random addition of nodes, showing the benefits of the adaptive construction employed in RADIS.

RADIS requires N evaluations of the posterior $\pi(\mathbf{x})$ at each iteration, so that the total number of posterior evaluations is $E = N_0 + NT$. Let denote as $C_{\text{eval-post}}$ the cost of evaluating $\pi(\mathbf{x})$ once, so that the total cost of evaluating the posterior is $EC_{\text{eval-post}}$. In addition to E posterior evaluations, RADIS carries out different other tasks, namely (i) evaluate L times the current emulator per iteration, (ii) perform N resampling steps per iteration over L possible samples, and (iii) compute the denominator of the final IS weights at the end of the algorithm. Let $C_{\text{eval-emulator}}$, $C_{\text{resampling}}$ and $C_{\text{den-weights}}$ denote the total costs after T iterations of RADIS, associated to tasks (i)-(iii). In term of computational time, RADIS can be convenient with respect to other schemes, when the inequality

$$C_{\text{eval-post}} > \frac{1}{E} (C_{\text{eval-emulator}} + C_{\text{resampling}} + C_{\text{den-weights}}), \quad (22)$$

is fulfilled. For an example, see the numerical experiment in Section 8.3 and the results in Table 9. Recall that all the values $C_{\text{eval-post}}$, $C_{\text{eval-emulator}}$, $C_{\text{resampling}}$, and $C_{\text{den-weights}}$ also depend on the specific implementation and language of the code and the different processors/machines.

Generally, the term $C_{\text{eval-emulator}}$ dominates the other two since it is composed of evaluating L times the emulator for T iterations. Moreover, due to the non-parametric construction and the fact that we increase the set of active nodes in N , evaluating the interpolator becomes more costly with the iterations. More specifically, in the NN based approach, after T iterations we have $C_{\text{eval-emulator}} \approx \sum_{t=1}^T \mathcal{O}(Lnt) = \mathcal{O}(LNT^2)$. In the GP-AIS scheme, we have the additional cost of inverting the $J_t \times J_t$ matrix at each iteration (recall that $J_t = N_0 + N(t-1)$). This cost at each iteration is $\mathcal{O}(J_t^3) \approx \mathcal{O}(N^3 t^3)$, for t big enough. Then, in GP-AIS, $C_{\text{eval-emulator}} \approx \sum_{t=1}^T \mathcal{O}(N^3 t^3) + \mathcal{O}(LNT^2) = \mathcal{O}(N^3 T^4) + \mathcal{O}(LNT^2)$.

Table 2

Parsimonious update in Step 3 of Table 1.

<p>- Initialization: Choose an acceptance function $p_A(\mathbf{x})$, set $\mathcal{S}_t = \mathcal{S}_{t-1}$, and consider the cloud of resampled particles $\{\mathbf{x}_{t,n}\}_{n=1}^N$, from the previous step of Table 1.</p> <p>- For $n = 1, \dots, N$:</p> <ol style="list-style-type: none"> 1. Draw $u \sim \mathcal{U}([0, 1])$. 2. If $u \leq p_A(\mathbf{x}_{t,n})$, then set $\mathcal{S}_t = \mathcal{S}_t \cup \{\mathbf{x}_{t,n}\}$. Otherwise, if $u > p_A(\mathbf{x}_{t,n})$, discard $\mathbf{x}_{t,n}$. <p>- Output: Return \mathcal{S}_t and $J_t = \mathcal{S}_t$.</p>

Table 3

Alternative parsimonious update considering a sequential updating of the emulator.

<p>- Initialization: Set $\hat{\pi}_t^{(0)}(\mathbf{x}) = \hat{\pi}_t(\mathbf{x})$, choose an acceptance function $p_A^{(0)}(\mathbf{x}) = p_A(\mathbf{x} \hat{\pi}_t^{(0)})$, set $k = 0$ and $\mathcal{S}_t = \mathcal{S}_{t-1}$, and consider the cloud of resampled particles $\{\mathbf{x}_{t,n}\}_{n=1}^N$, from the previous step of Table 1. Note that, more generally, $p_A^{(k)}(\mathbf{x}) = p_A(\mathbf{x} \hat{\pi}_t^{(k)})$ where $k \geq 0$ is an index.</p> <p>- For $n = 1, \dots, N$:</p> <ol style="list-style-type: none"> 1. Draw $u \sim \mathcal{U}([0, 1])$. 2. If $u \leq p_A^{(k)}(\mathbf{x}_{t,n})$, then set $\mathcal{S}_t = \mathcal{S}_t \cup \{\mathbf{x}_{t,n}\}$, and update the emulator construction $\hat{\pi}_t^{(k+1)}(\mathbf{x})$ considering the new set \mathcal{S}_t. Set also $k \leftarrow k + 1$. <p>- Output: Return \mathcal{S}_t, $J_t = \mathcal{S}_t$ and $\hat{\pi}_{t+1}(\mathbf{x}) = \hat{\pi}_{t+1}(\mathbf{x}; \mathcal{S}_t) = \hat{\pi}_t^{(k)}(\mathbf{x})$.</p>
--

In the next section, we describe different procedures to decrease $C_{\text{eval-emulator}}$.

6. Construction of parsimonious emulators

So far, we have considered updating the interpolant at each iteration t by adding all the N samples drawn at that iteration. In order to control the computational cost of evaluating the emulator, we can design a strategy for accepting or rejecting some of the possible additional nodes. This can be done assigning acceptance probabilities, $p_A(\mathbf{x}_{t,n}) \in [0, 1]$, to each of the N samples (in the same fashion of [27,64]). Therefore, the update part of Step 3 in Table 1 would be replaced by the routine in Table 2.

Proper acceptance functions We say that an acceptance probability, $p_A(\mathbf{x}) : \mathcal{X} \rightarrow [0, 1]$, is *proper* if it satisfies

$$\mathbf{C1:} \quad p_A(\mathbf{x}) \rightarrow 0, \quad \text{if} \quad |\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})| \rightarrow 0, \quad (23)$$

for any $\mathbf{x} \in \mathcal{X}$, and

$$\mathbf{C2:} \quad p_A(\mathbf{x}) = 0 \text{ if and only if } |\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})| = 0. \quad (24)$$

Hence, for any node contained already in \mathcal{S}_{t-1} , i.e., $\mathbf{z} \in \mathcal{S}_{t-1}$, we have $p_A(\mathbf{z}) = 0$. For this reason, as we show below, the acceptance function often depends on the current emulator $\hat{\pi}_t(\mathbf{x})$, i.e., we should write $p_A(\mathbf{x}) = p_A(\mathbf{x}|\hat{\pi}_t)$. Hence, a more precise and parsimonious construction would consider a sequential updating of the emulator (since $p_A(\mathbf{x})$ also should change during the acceptance tests), as shown in Table 3.

Remark 12. Note that the procedures in Table 2 and 3 do not require additional evaluations of the target π and the interpolator $\hat{\pi}_t$, since all the values $\pi(\mathbf{x}_{t,n})$ and $\hat{\pi}_t(\mathbf{x}_{t,n})$, for all n , are already obtained.

The difference between the schemes in Tables 2 and 3, in term of performance and computational cost, becomes more relevant as N grows. Note that the order of the tests in Table 3 could be also

relevant and some strategies for ordering $\{\mathbf{x}_{t,n}\}_{n=1}^N$ (in a suitable way) could be designed. Below, we introduce some examples of proper acceptance functions and also some reasonable improper ones.

6.1. Examples of proper acceptance functions

One possibility of proper acceptance function is

$$\mathbf{A1:} \quad p_A(\mathbf{x}) = 1 - \frac{\min\{\pi(\mathbf{x}), \hat{\pi}_t(\mathbf{x})\}}{\max\{\pi(\mathbf{x}), \hat{\pi}_t(\mathbf{x})\}} = \frac{|\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})|}{\max\{\pi(\mathbf{x}), \hat{\pi}_t(\mathbf{x})\}}, \quad (25)$$

where we have used $|\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})| = \max\{\pi(\mathbf{x}), \hat{\pi}_t(\mathbf{x})\} - \min\{\pi(\mathbf{x}), \hat{\pi}_t(\mathbf{x})\}$. Another possibility is to consider both the discrepancy between π and $\hat{\pi}_t$, and the distance to the closest node $\mathbf{s}^* \in \mathcal{S}_{t-1}$ to \mathbf{x} , i.e.,

$$\mathbf{A2:} \quad p_A(\mathbf{x}) = \left(1 - e^{-\alpha|\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})|}\right) \left(1 - e^{-\beta\|\mathbf{x} - \mathbf{s}^*\|}\right), \quad \alpha, \beta \geq 0. \quad (26)$$

If either $\alpha = 0$ or $\beta = 0$ (or both), then $p_A(\mathbf{x}) = 0$. As $\alpha \rightarrow \infty$ and $\beta \rightarrow \infty$ grow, then $p_A(\mathbf{x}) \rightarrow 1$. When $\alpha = \infty$ and β is finite, then $p_A(\mathbf{x}) = 1 - e^{-\beta\|\mathbf{x} - \mathbf{s}^*\|}$ and the acceptance probability is bigger when the point \mathbf{x} is far from its closest node, i.e., we have a space-filling strategy. When α is finite and $\beta \rightarrow \infty$, then $p_A(\mathbf{x}) = 1 - e^{-\alpha|\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})|}$, and the acceptance probability is bigger if there is a large discrepancy between π and the interpolant at \mathbf{x} . Thus, unlike in (25), in (26) we should tune the values α , and β according to the computational budget we have, or according to the trade-off between computational effort and performance.

Note that, in the acceptance functions above, we have $p_A(\mathbf{x}) \in [0, 1]$ for all \mathbf{x} , and the condition (23) is fulfilled. Moreover, these acceptance functions depend only on \mathbf{x} , $\pi(\mathbf{x})$ and $\hat{\pi}_t(\mathbf{x})$. The decision is done considering the quality of the approximation of $\hat{\pi}_t(\mathbf{x})$ and, in Eq. (26), the relative position of \mathbf{x} with respect to the nodes in \mathcal{S}_{t-1} . They do not depend on the rest of $N - 1$ possible nodes within $\{\mathbf{x}_{t,n}\}_{n=1}^N$ to be tested. Nevertheless, if we use the sequential updating scheme of Table 3, the acceptance probability will change depending on the order in which we test the candidate nodes.

An example of proper acceptance function depending on the population of candidate nodes is described next. Let us define $R(\mathbf{x}) = |\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})|$. Considering $\mathbf{x} \in \{\mathbf{x}_{t,1}, \dots, \mathbf{x}_{t,N}\}$ (i.e., one point within the set of possible nodes to be included) and defining $R_{\max} = \max_{\mathbf{x} \in \{\mathbf{x}_{t,1}, \dots, \mathbf{x}_{t,N}\}} R(\mathbf{x})$, we can set

$$\mathbf{A3:} \quad p_A(\mathbf{x}) = \frac{R(\mathbf{x})}{R_{\max}}, \quad \text{with } \mathbf{x} \in \{\mathbf{x}_{t,n}\}_{n=1}^N. \quad (27)$$

Again $p_A(\mathbf{x}) \in [0, 1]$ and the condition (23) is satisfied. Note that a normalization of $R(\mathbf{x})$ using $\sum_{n=1}^N R(\mathbf{x}_{t,n})$ instead of R_{\max} would produce very small acceptance probabilities as N grows (note that $R(\mathbf{x}) \geq 0$ for all \mathbf{x}). This is a non beneficial effect in our opinion, since the decrease of $p_A(\mathbf{x})$ is not due to a good quality of the approximation $\hat{\pi}_t$, but is generated by the increase of the possible alternative denominator $\sum_{n=1}^N R(\mathbf{x}_{t,n})$. Resampling schemes could be also employed but provide improper acceptance functions, as we discuss below.

6.2. Examples of improper acceptance functions

Let us define the auxiliary weights $\rho(\mathbf{x}) = \frac{F(\mathbf{x})}{\hat{\pi}_t(\mathbf{x})}$ where $F(\mathbf{x})$ is function that can chosen in different ways, $F(\mathbf{x}) = \pi(\mathbf{x})$, $F(\mathbf{x}) = |\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})|$ or $F(\mathbf{x}) = |\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})|\hat{\pi}_t(\mathbf{x})$, for instance. The nodes to be included are then selected resampling N times within the set $\{\mathbf{x}_{t,n}\}_{n=1}^N$ according to the following probability mass,

Table 4

Parsimonious update in Step 3 of Table 1 based on resampling.

- **Initialization:** Choose a numerator function $F(\mathbf{x})$ (e.g., $F(\mathbf{x}) = \pi(\mathbf{x})$ or $F(\mathbf{x}) = |\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})|$) for the weight $\rho(\mathbf{x}) = \frac{F(\mathbf{x})}{\hat{\pi}_t(\mathbf{x})}$. Set $\mathcal{S}_t = \mathcal{S}_{t-1}$, and consider the cloud of resampled particles $\{\mathbf{x}_{t,n}\}_{n=1}^N$, from the previous step of Table 1. Then:

1. Resample N times within $\{\mathbf{x}_{t,n}\}_{n=1}^N$ according to the probability mass defined as

$$\bar{\rho}_{t,i} = \bar{\rho}(\mathbf{x}_{t,i}) = \frac{\rho(\mathbf{x}_{t,i})}{\sum_{n=1}^N \rho(\mathbf{x}_{t,n})}, \quad i = 1, \dots, N,$$

obtaining the new set $\{\tilde{\mathbf{x}}_{t,n}\}_{n=1}^N$.

2. Take the unique values in $\{\tilde{\mathbf{x}}_{t,n}\}_{n=1}^N$ (i.e., removing the repetitions) obtaining $\{\mathbf{v}_{t,k}\}_{k=1}^K$ (where K is the number of unique values in $\{\tilde{\mathbf{x}}_{t,n}\}_{n=1}^N$).
3. Set $\mathcal{S}_t = \mathcal{S}_{t-1} \cup \{\mathbf{v}_{t,1}, \dots, \mathbf{v}_{t,K}\}$.

-**Output:** Return \mathcal{S}_t and $J_t = |\mathcal{S}_t|$.

$$\bar{\rho}(\mathbf{x}_{t,i}) = \frac{\rho(\mathbf{x}_{t,i})}{\sum_{n=1}^N \rho(\mathbf{x}_{t,n})}, \quad i = 1, \dots, N,$$

and taking only the *unique* values (i.e., without repetitions). Table 4 summarizes this idea.

The acceptance probability is, in this case,

$$p_A(\mathbf{x}_{t,i}) = 1 - (1 - \bar{\rho}(\mathbf{x}_{t,i}))^N. \quad (28)$$

Thus, the procedure in Table 4 is equivalent (in term of number of added nodes) to apply the procedure in Table 2 and $p_A(\mathbf{x})$ in (28) above. Observe also that, with these schemes, even in the ideal case $\hat{\pi}_t(\mathbf{x}) = \pi(\mathbf{x})$ for all \mathbf{x} , we always add at least one node to the new sets \mathcal{S}_t (i.e., $K \geq 1$). This is due to the *improperness* of the acceptance functions. Then, these resampling-based schemes could possibly yield less parsimonious emulators. Nevertheless, they are easy to implement. Starting from the samples $\mathbf{z}_{t,\ell} \sim q_{\text{aux}}(\mathbf{x})$ in RADIS, the added points $\{\mathbf{v}_{t,k}\}_{k=1}^K$ in Table 4 are then obtained as results of two resampling procedures and finally considering the unique values:

$$\{\mathbf{z}_{t,\ell}\}_{\ell=1}^L \xrightarrow{\bar{\gamma}_{t,\ell}} \{\mathbf{x}_{t,n}\}_{n=1}^N \xrightarrow{\bar{\rho}_{t,\ell}} \{\mathbf{v}_{t,k}\}_{k=1}^K.$$

In the vanilla version of RADIS, the nodes are obtained applying just the first resampling at each iteration. Another example of *improper* acceptance function that is not based on a resampling procedure (and does not take into account all the population $\{\mathbf{x}_{t,n}\}_{n=1}^N$, jointly) is

$$p_A(\mathbf{x}) = \begin{cases} 1 & \text{if } |\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})| > \epsilon, \\ 0 & \text{if } |\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})| \leq \epsilon, \end{cases} \quad \text{for } \epsilon \geq 0. \quad (29)$$

Note that for a finite positive value of $\epsilon > 0$, after some iterations, possibly we will have $p_A(\mathbf{x}) = 0$, i.e., the adaptation of the emulator is stopped. This is the reason of its *improperness*, since it does not fulfill C2. If $\epsilon = 0$, then we always have $p_A(\mathbf{x}) = 1$, adding all the nodes. If $\epsilon = \infty$, we have always $p_A(\mathbf{x}) = 0$, and we never update the emulator. With a suitable choice of ϵ (tuned according to computational budget available), this acceptance function can be also a good option. A numerical comparison among these acceptance probabilities is given in Section 8.

7. RADIS for model emulation and sequential inversion

In this section, we describe the application of RADIS to solve Bayesian inverse problems. We have already considered the case of obtaining a surrogate function for the (unnormalized) density π (or $\log \pi$). We here focus on inverse inference problems where our aim is also to obtain an emulator of the costly forward model.

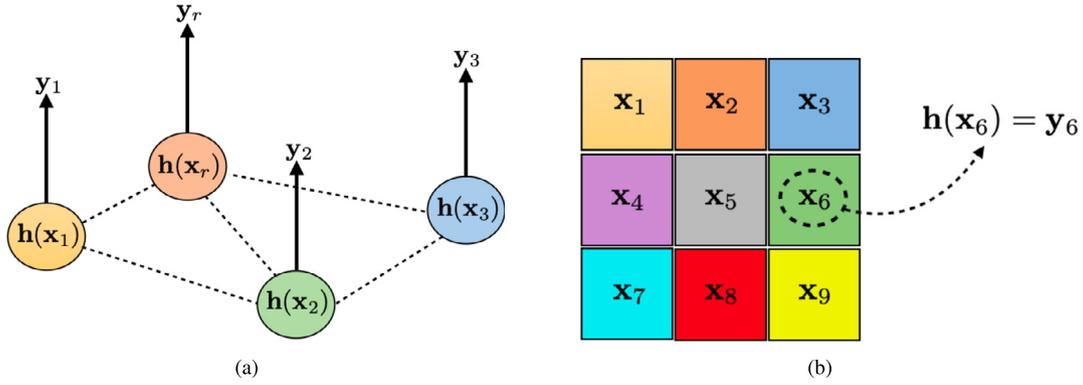


Fig. 4. (a) Different inversion problems related to each other involving the same underlying physical model $\mathbf{h}(\cdot)$. Their relationships are represented by (dashed lines) edges between the nodes. (b) Example of network in an image, where each pixel represents a node of the network. This is the scenario in remote sensing image processing, where \mathbf{x}_i represents the physical state parameters to infer from a set of acquired (or simulated) spectra \mathbf{y}_i (in this figure, we consider noise-free observations).

More specifically, let us consider a generic Bayesian inversion problem

$$\mathbf{y} = \mathbf{h}(\mathbf{x}) + \mathbf{v}, \quad (30)$$

where $\mathbf{h}(\mathbf{x}) : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_y}$ represents a non-linear mapping defining a physical or mechanistic model (e.g. a complex energy transfer model, a climate model subcomponent integrating subgrid physical processes, or a set of differential equations describing a chemical diffusion process) and \mathbf{v} has a multivariate Gaussian pdf (e.g., with zero mean and a diagonal covariance matrix with σ^2 in the diagonal). Considering a prior $g(\mathbf{x})$ over \mathbf{x} , the posterior is

$$\bar{\pi}(\mathbf{x}) \propto \pi(\mathbf{x}) = \exp\left(-\frac{1}{2\sigma^2}\|\mathbf{y} - \mathbf{h}(\mathbf{x})\|^2\right) g(\mathbf{x}),$$

which can be costly to evaluate if $\mathbf{h}(\mathbf{x})$ is a complex model. In this setting, it is often required to build an emulator of the physical model $\mathbf{h}(\mathbf{x})$ instead of a surrogate function for the pdf π [65–67,14]. However, we can build $\hat{\mathbf{h}}(\mathbf{x})$ using the same procedures in Sect. 4.2, and then obtain

$$\hat{\pi}(\mathbf{x}) = \exp\left(-\frac{1}{2\sigma^2}\|\mathbf{y} - \hat{\mathbf{h}}(\mathbf{x})\|^2\right) g(\mathbf{x}),$$

which can be employed as proposal in our scheme. Hence, in this case, we obtain two emulators: $\hat{\mathbf{h}}(\mathbf{x})$ of the physical model, and $\hat{\pi}(\mathbf{x})$ of the posterior.

In many real-world applications, we have a sequence of inverse problems

$$\mathbf{y}_r = \mathbf{h}(\mathbf{x}_r) + \mathbf{v}_r, \quad r = 1, \dots, R, \quad (31)$$

where R denotes the number of observation nodes in the network, but the physical model \mathbf{h} is the same for all nodes. See an illustrative example in Fig. 4(a). The underlying graph represents different features and may have different statistical meanings. Moreover, it can contain prior information directly given in the specific problem. As an example, consider the case of an image where each pixel is represented as a node in the network, see Fig. 4(b), and the goal is to retrieve a set of parameters \mathbf{x} from the observed or simulated pixels \mathbf{y} . This is the standard scenario in remote sensing applications, where the observations \mathbf{y} are very high dimensional (depending on the sensory system and satellite platform ranging from a few spectral channels to even thousands) and the set of parameters \mathbf{x} describe the physical characteristics of each particular observation (e.g. leaf or canopy structure, observation characteristics, vegetation health and status, etc.). In other settings the graph

must be also inferred, i.e., the connections should be learned as well. A simple strategy is to consider the strength of the link is proportional to $\exp(-\|\mathbf{y}_r - \mathbf{y}_j\|)$, for instance. Other more sophisticated procedures can be also employed [68]. Given Eq. (31), a piece of the likelihood function is

$$p(\mathbf{y}_r | \mathbf{x}_r) \propto \exp\left(-\frac{1}{2\sigma^2}\|\mathbf{y}_r - \mathbf{h}(\mathbf{x}_r)\|^2\right), \quad r = 1, \dots, R.$$

Note that the observation model $\mathbf{h}(\cdot)$ is shared in all the R nodes. The complete likelihood function is $p(\mathbf{y}_{1:r} | \mathbf{x}_{1:r}) = p(\mathbf{y}_1, \dots, \mathbf{y}_R | \mathbf{x}_1, \dots, \mathbf{x}_R) = \prod_{i=1}^R p(\mathbf{y}_i | \mathbf{x}_i)$. A complete Bayesian analysis can be considered in this scenario, implementing also RADIS within a particle filter for an efficient inference. However, it is out of the scope of this work and we leave it as a future research line.

8. Numerical experiments

In this section, we provide several numerical tests in order to show the performance of the proposed scheme and compare them with benchmark approaches in the literature. The first example corresponds to a nonlinear banana shaped density in dimension $d_x = 2$, where we compare NN-AIS against standard IS algorithms. The second test is a multimodal scenario with dimension $d_x = 10$, where we test the combination of an AIS algorithm with NN-AIS against other AIS. An application to an astronomical model is also given, where we provide a comparison in terms of computation time. Finally, we consider an application to remote sensing, specifically, we test our scheme in multiple Bayesian inversions of PRO-SAIL.

8.1. Toy example 1: banana-shaped density

We consider a banana shaped target pdf,

$$\bar{\pi}(\mathbf{x}) \propto \exp\left(-\frac{(\eta_1 - Bx_1 - x_2^2)^2}{2\eta_0^2} - \sum_{i=1}^{d_x} \frac{x_i^2}{2\eta_i^2}\right), \quad (32)$$

with $B = 4$, $\eta_0 = 4$ and $\eta_i = 3.5$ for $i = 1, \dots, d_x$, where $\mathcal{X} = [-10, 10] \times [-10, 10]$, i.e., bounded domain. We consider $d_x = 2$ and compute in advance Z and the mean of the target (i.e., the groundtruth) by using a costly grid, so that we can check the performance of the different techniques.

8.1.1. Estimating Z and μ

We aim to estimate $Z = 7.9976$ and $\mu = [-0.4841, 0]$ with NN-AIS and compare it, in terms of relative mean squared error (RMSE), with different IS algorithms considering the same number

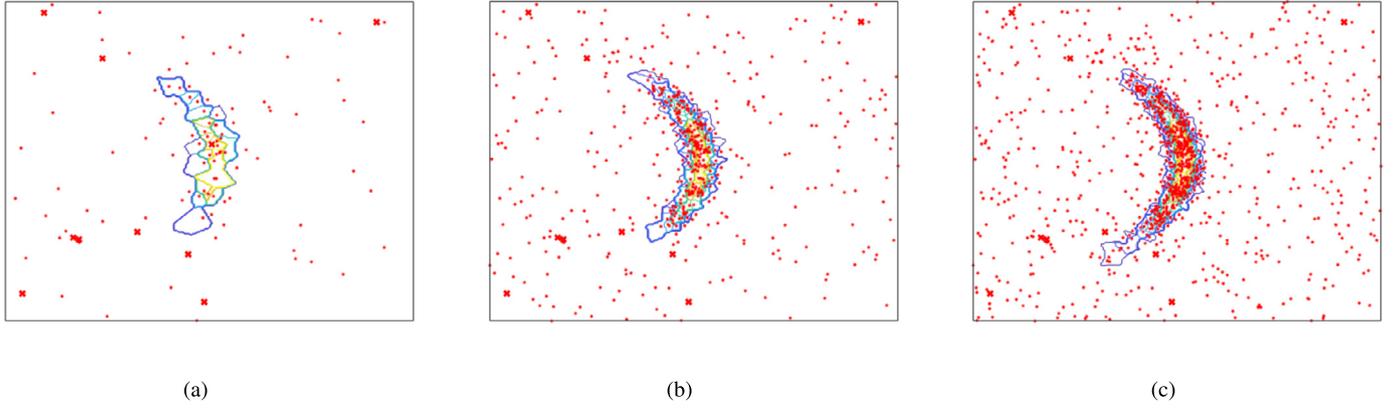


Fig. 5. Evolution of $\hat{\pi}_t$ from NN-AIS+U through iterations (a) $t = 10$, (b) $t = 50$, and (c) $t = 100$.

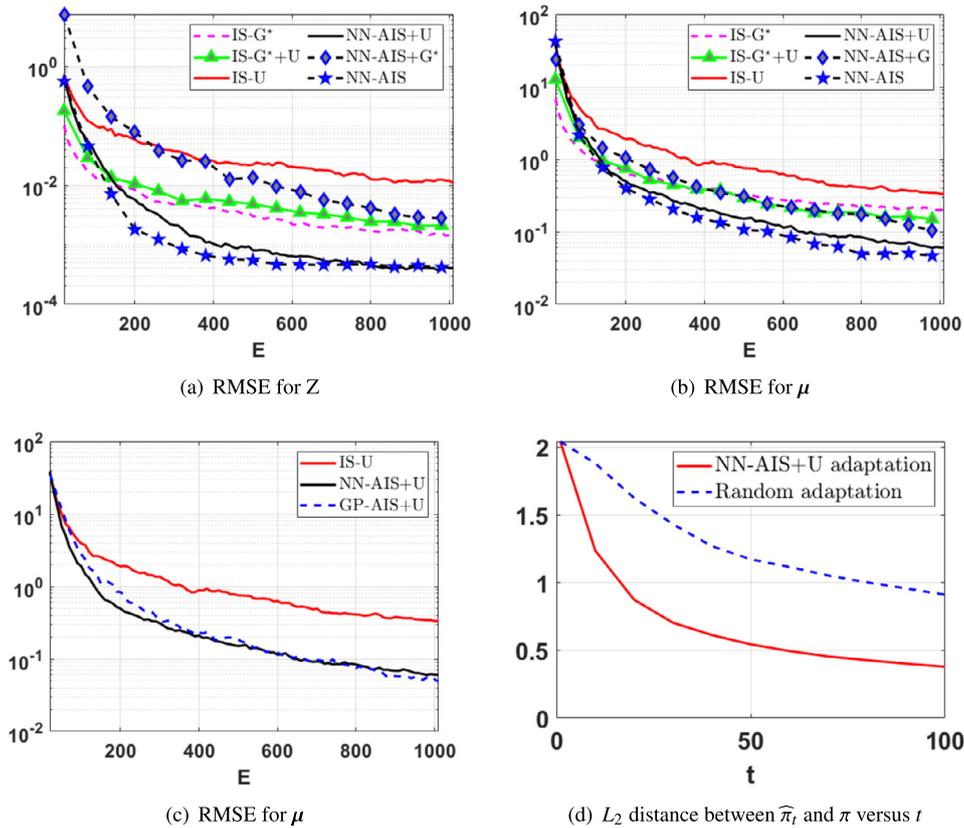


Fig. 6. (a) RMSE in log-scale for Z as function of evaluations E . (b) RMSE in log-scale for μ as function of E . (c) RMSE of GP-AIS+U in log-scale for μ as function of E . (d) L_2 distance between π and $\hat{\pi}_t$ when the nodes are adaptively obtained by NN-AIS+U (in solid line), and when the nodes are random and uniformly chosen in the domain (in dashed line), as a function of t .

of target evaluations. The results are averaged over 500 independent simulations. The goal is to investigate the performance of NN-AIS as compared to other parametric IS algorithms that consider a proposal, well designed in advance. We set $T = 100$ and $N = 10$, and use 10 starting nodes (random chosen in the domain) to build $\hat{\pi}_1(\mathbf{x}|\mathcal{S}_0)$. With the selected values of T and N the total budget of target evaluations is $E = 10 + NT = 1010$.

Methods We consider three variants of NN-AIS to illustrate three different scenarios: in the first one (denoted as **NN-AIS**) initial nodes uniform in $[-10, 10] \times [-10, 10]$, i.e. good initialization, without $\bar{q}_{\text{par}}(\mathbf{x})$; (**NN-AIS+U**) same initialization with $\bar{q}_{\text{par}}(\mathbf{x}) = \frac{1}{|\mathcal{X}|}$, i.e. good initialization and with a good choice of $\bar{q}_{\text{par}}(\mathbf{x})$; (**NN-AIS+G**) initial nodes are uniform in $[5, 10] \times [5, 10]$ with Gaussian

$\bar{q}_{\text{par}}(\mathbf{x}) = \mathcal{N}(\mathbf{x}|[2, 2]^\top, 3^2\mathbf{I}_2)$, i.e., a bad initialization with a bad choice of the parametric proposal $\bar{q}_{\text{par}}(\mathbf{x})$. In all cases, we consider a fixed value of $\alpha_t = \frac{1}{2}$.

Furthermore, we compare the NN-AIS schemes with three alternative IS methods: (**IS-U**) with uniform proposal in \mathcal{X} , which is very good choice of proposal in this problem; (**IS-G***) with Gaussian proposal matching the moments of $\bar{\pi}(\mathbf{x})$, i.e., the optimal Gaussian proposal; (**IS-G*+U**) with a proposal which is an equally weighted mixture of the two previous cases. In addition, we also test our algorithm using GPs, denoted GP-AIS+U.

Discussion As shown in Figs. 6(a)-(b), NN-AIS and NN-AIS+U outperform the rest. NN-AIS performs a bit better than NN-AIS+U: the use of a parametric proposal is safer but entails a loss of

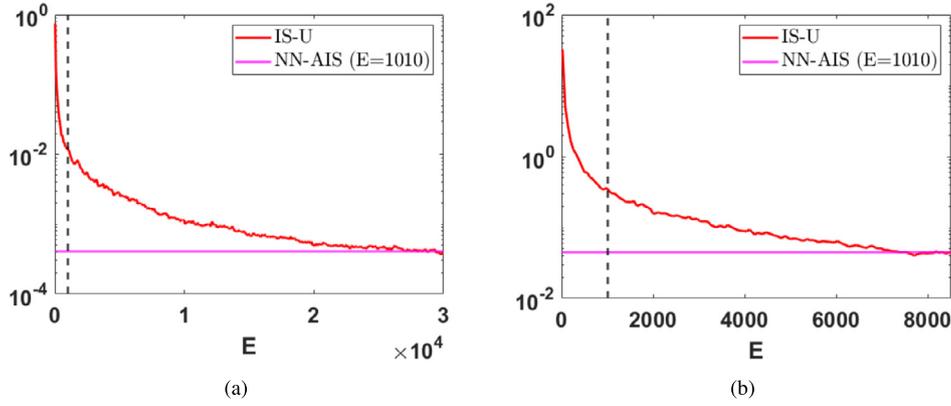


Fig. 7. We show the number of additional evaluations required by IS-U to achieve the same RMSE than NN-AIS with $E = 1010$ in (a) the estimation of Z , and (b) the estimation of μ . The red line represents the RMSE of IS-U as a function of E , while the horizontal line is the RMSE achieved by NN-AIS with $E = 1010$. The vertical dash line is at $E = 1010$. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

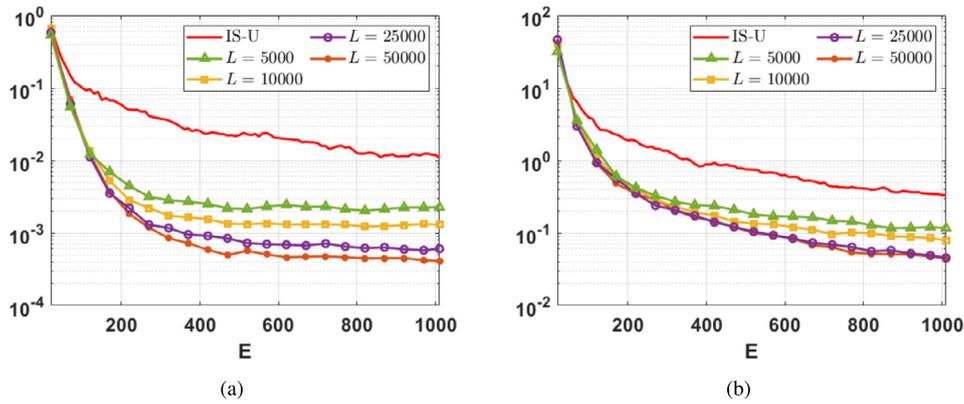


Fig. 8. Performance of NN-AIS+U with different choices of $L \in \{1000, 5000, 10000, 25000, 50000\}$ in (a) the estimation of Z , and (b) the estimation of μ . The red curve represents the RMSE of IS-U as a function of E .

performance, trading off exploitation for exploration. In Fig. 6(a), NN-AIS+G shows worse performance in estimating Z in the early iterations as a consequence of the bad initialization and bad parametric proposal. However, it quickly improves and start performing as good as IS-G* and IS-G*+U. In Fig. 6(b), regarding the estimation of μ , our methods perform better than alternative IS algorithms. Fig. 6(c) shows that GP-AIS+U provides similar performance than NN-AIS+U. Overall, this simple experiment shows the range of performance of our method: it is best if we use only our method, provided that we have a good initialization; adding a good parametric proposal is safer if we do not trust our initialization, showing just a small loss of performance w.r.t. the first scenario. In the case both the initialization and parametric proposal are wrongly chosen, our method is able to achieve good results and recover quickly from a bad initialization.

Additional comparison We have run IS-U for $E > 1010$ until it reached the same error in estimation achieved by NN-AIS. The results are depicted in Figs. 7. Specifically, in Fig. 7(a) we see that around 29000 more evaluations are needed to obtain the same error in estimating Z , and Fig. 7(b) shows that around 7000 more evaluations to obtain the same error in estimating μ .

8.1.2. Convergence of $\hat{\pi}_t$ to π

The convergence of $\hat{\pi}_t$ to π depends on the fact that nodes should fill the space enough (see Appendix A). However, some filling strategies yield a faster convergence than others. In our simulations, we aim to show that the construction provided by NN-AIS+U converges faster than another construction using nodes

random and uniformly chosen in the domain \mathcal{X} . Figs. 5 and 6(d) show that the approximation $\hat{\pi}_t$ obtained by NN-AIS+U is indeed converging to π as t increases. In Fig. 6(c), we show the L_2 distance between π and $\hat{\pi}_t$ with random nodes (in dashed line), and by NN-AIS+U (in solid line), along with the number of iterations t . As shown in Fig. 6(d), the $\hat{\pi}_t$ gets more rapidly closer to π in L_2 when the nodes are sampled from NN-AIS+U rather than only adding random points, uniformly over the domain.

8.1.3. Comparing NN-AIS+U with different values of L

In our proposed approach, we need to evaluate L times the approximation $\hat{\pi}_t$ at each iteration. The computation cost of the algorithm thus scales with L , which needs to be big enough (and bigger than N) so that the resampling step and the estimation of c_t are accurate. Here, we investigate the performance of NN-AIS+U for several values $L \in \{5000, 10000, 25000, 50000\}$. As expected, Fig. 8 shows that the performance of the algorithm deteriorates as we lower the value of L . However, note that all NN-AIS scheme with the considered L perform better compared to standard IS with uniform proposal.

8.1.4. Results of the parsimonious constructions

In the vanilla version of RADIS, the approximation $\hat{\pi}_t$ is refined by adding the N samples drawn at iteration t to the set of active nodes. Since we consider non-parametric approximations, this implies that $\hat{\pi}_t$ becomes more complex, i.e. more costly to evaluate, as t grows. In Sect. 6, we showed means of controlling the complexity of $\hat{\pi}_t$ by the computation of acceptance probabilities: instead of adding all the samples, the n -th sample is added with

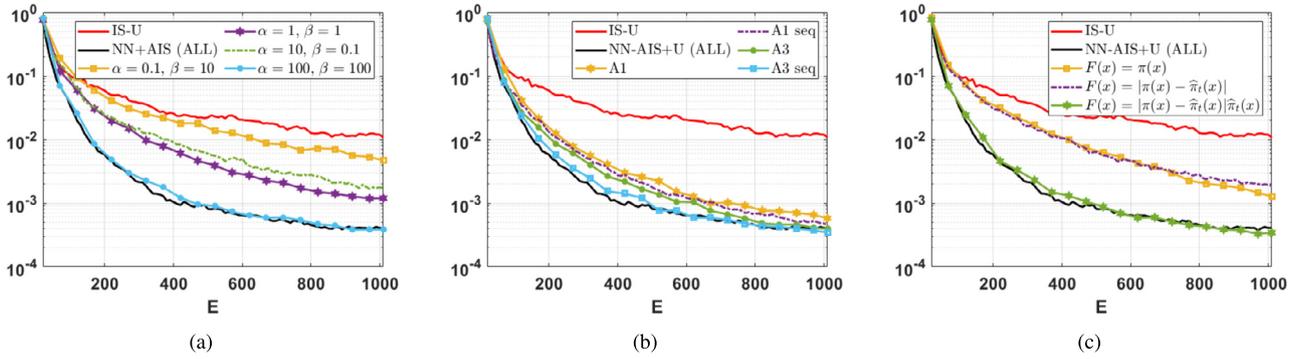


Fig. 9. Performance of NN-AIS+U with acceptance function from Eq. (26) for different choices of α and β , in (a) the estimation of Z . Number of nodes versus t in (b) in linear scale, and (c) in logarithm scale. The black solid curve represents the number of nodes of NN-AIS+U that accepts all.

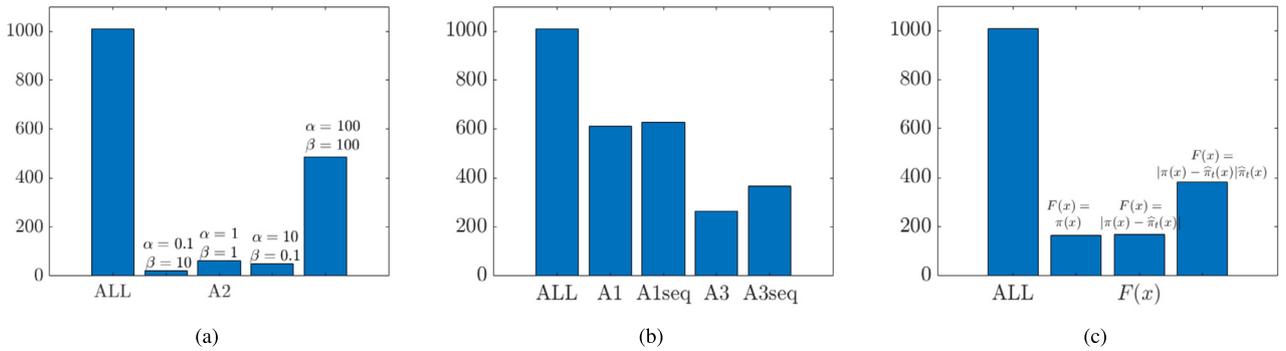


Fig. 10. Final number of added nodes for the construction of the emulator for NN-AIS+U with several acceptance functions of different parsimonious schemes. Figure (a) is devoted to the strategy A2 with different parameters, Figure (b) to the strategies A1 and A3, and Figure (c) the resampling strategies in Table 4.

certain probability. Here, we test the application of several acceptance probabilities to NN-AIS+U and compare the performance with respect to NN-AIS+U that accepts all nodes. We also examine the complexity, in terms of number of nodes, of the final emulator. Specifically, we consider the acceptance functions **A1** in Eq. (25), **A2** in Eq. (26) and **A3** in Eq. (27). We also test three variants of the improper acceptance function in Sect. 6.2, namely $F(\mathbf{x}) = \pi(\mathbf{x})$, $F(\mathbf{x}) = |\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})|$ and $F(\mathbf{x}) = |\pi(\mathbf{x}) - \hat{\pi}_t(\mathbf{x})|\hat{\pi}_t(\mathbf{x})$. The results are given in Figs. 9, Fig. 10 and Fig. 11. Note that NN-AIS+U (ALL) represents the vanilla version NN-AIS+U in Table 1, adding all the nodes at the Step 3.

Fig. 9(a) shows the application of the acceptance probability **A2** for different choices of α and β using the updating scheme in Table 3. Recall that, when α or β are 0, the acceptance probability is 0. When $\alpha \gg 1$ and $\beta > 0$, the nodes are added in a space-filling fashion. On the contrary, when $\beta \gg 1$ and $\alpha > 0$, the nodes are added by accounting for the discrepancy between π and $\hat{\pi}_t$. We note that the former strategy works better than the latter, as shown in Fig. 9(a). Moreover, the performance is better when $\alpha = \beta = 1$, that is, both strategies at the same time. As α and β grow, we recover the performance of the NN-AIS+U accepting all samples. Fig. 10(a) shows the number of nodes of the final constructed emulators. We see that the choice $\alpha = \beta = 100$ produces an approximation $\hat{\pi}_t$ that has only half of the nodes of the algorithm accepting all the samples, but achieves the same level of precision in the estimation. We also tested the acceptance functions based on resampling in Eq. (28). The results are given in Figs. 9(c) and 10(c). We also tested the acceptance functions **A1** and **A3**, each one with the two possible updating schemes from Tables 2 (non-sequential) and 3 (sequential). As shown in Fig. 9(b), the acceptance function **A3** provides better results than **A1**. For both, the use of a sequential updating scheme improve the results. Fig. 10(b) shows the number of final nodes of the emulator. We

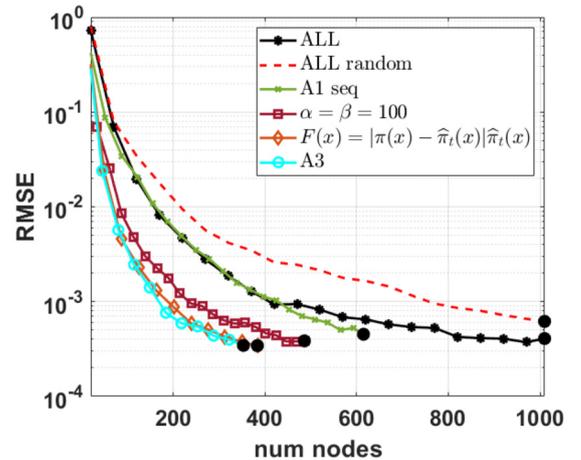


Fig. 11. RMSE for NN-AIS+U with different acceptance functions (choosing the best schemes in the previous tests), versus the total number of added nodes at each iteration. We have also incorporated a curve (depicted with dashed line) of an NN-AIS+U scheme where N nodes are added at each iteration but chosen randomly in the space (instead of adding the nodes obtained in the inner resampling in Step 3 of Table 1). The end point in each curve is highlighted with greater black circle. The reason is that this last point is completely comparable among the different curve since, at this point, we have the same number of target evaluations E . Observing these end points, we see that all the parsimonious schemes shown in the figure provide the same or smaller error than the vanilla NN-AIS+U, with a smaller number of added nodes.

can observe that several parsimonious schemes provide very good performance, close to the vanilla NN-AIS+U (with a much smaller number of added nodes).

Finally, in Fig. 11 we compare the best parsimonious schemes with the vanilla NN-AIS+U method, showing their RMSE as function of the total number of added nodes at each iterations. Fur-

Table 5
MAE for Z with $E = 1000$ (best and worst MAE of each method are boldfaced).

Methods		$\xi = 1$	$\xi = 2$	$\xi = 3$	$\xi = 4$	$\xi = 5$	$\xi = 6$
PMC	$N_{\text{PMC}} = 10$	0.9993	0.9526	0.8603	0.6743	0.6024	0.6155
	$N_{\text{PMC}} = 100$	0.9998	0.9896	0.8853	0.6761	0.5192	0.4544
	$N_{\text{PMC}} = 200$	1.0002	0.9893	0.8816	0.7099	0.6389	0.5384
	$N_{\text{PMC}} = 500$	0.9995	0.9916	0.9741	0.8700	0.7421	0.6544
DM-PMC	$N_{\text{PMC}} = 10$	0.9991	0.9478	0.8505	0.6009	0.5352	0.5814
	$N_{\text{PMC}} = 100$	0.9997	0.8719	0.4490	0.2425	0.1901	0.2193
	$N_{\text{PMC}} = 200$	0.9999	0.9321	0.5708	0.3257	0.2374	0.2524
	$N_{\text{PMC}} = 500$	1.0000	0.9888	0.7969	0.5009	0.3684	0.3800
OC-LAIS	$N_{\text{LAIS}} = 10$	1.0000	1.0000	0.9992	0.9883	0.9468	0.9079
	$N_{\text{LAIS}} = 100$	0.9999	0.8731	0.4434	0.2785	0.2392	0.2870
	$N_{\text{LAIS}} = 200$	0.9982	0.7028	0.2418	0.1243	0.1406	0.2070
	$N_{\text{LAIS}} = 500$	0.9937	0.4949	0.1221	0.0857	0.1195	0.1786

thermore, as the dashed line in Fig. 6(d), we have compared with an NN-AIS+U scheme where N nodes are added at each iteration but chosen randomly in the space (instead of adding the nodes obtained in the inner resampling in Step 3 of Table 1). The corresponding curve is shown with a dashed line. The end point of each curve is highlighted with greater black circle. The reason is that this last point is completely comparable among the different curve since, at this point, we have the same number of target evaluations E . Therefore, observing these last points, we can see that all the parsimonious schemes achieve the same or smaller error than the vanilla NN-AIS+U, with a smaller number of added nodes. Employing the procedure **A3**, we save on average more than 40% of computational time with respect to the vanilla NN-AIS+U scheme, adding all the points.

8.2. Toy example 2: multimodal density

In this experiment, we consider a multimodal Gaussian target in $d_x = 10$,

$$\bar{\pi}(\mathbf{x}) = \frac{1}{3}\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1) + \frac{1}{3}\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2) + \frac{1}{3}\mathcal{N}(\mathbf{x}|\boldsymbol{\mu}_3, \boldsymbol{\Sigma}_3),$$

with $\boldsymbol{\mu}_1 = [5, 0, \dots, 0]$, $\boldsymbol{\mu}_2 = [-7, 0, \dots, 0]$, $\boldsymbol{\mu}_3 = [1, \dots, 1]$ and $\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_2 = \boldsymbol{\Sigma}_3 = 4^2\mathbf{I}_{10}$. We want to test the performance of the different methods in estimating the normalizing constant $Z = 1$. Specifically, our aim is to test the combination of our NN-AIS scheme with an AIS algorithm against other AIS algorithms. The budgeted of target evaluations is $E = 1000$.

Methods We consider three sophisticated AIS schemes, namely *population Monte Carlo* (PMC) [5], *layered adaptive IS* (LAIS) [41] and *adaptive multiple IS* (AMIS) [43]. These are AIS algorithms where the proposal (or proposals) gets updated at each iteration using information from previous samples. Specifically, PMC performs multinomial resampling to locate the proposals in the next iteration; AMIS matches the mean of the single proposal with the current estimation of the posterior mean using all previous samples; LAIS evolves the location parameters of the proposals with a MCMC algorithm. The goal is to compare the performance of PMC, LAIS and AMIS with a combination of our NN-AIS scheme and LAIS.

We set Gaussian pdfs as the proposal pdfs for all methods. We also need to set the number of these proposals in PMC and LAIS, as well as the dispersion of the Gaussian densities. For PMC, we test different number of proposals $N_{\text{PMC}} \in \{10, 100, 200, 500\}$, whose means are initialized at random in $[-15, 15]^{10}$. At each iteration of PMC, one sample is drawn from each of the N_{PMC} proposals, hence the algorithm is run for $T_{\text{PMC}} = \frac{1000}{N_{\text{PMC}}}$ iterations for a fair comparison. As a second alternative, we consider the deterministic mixture weighting approach for PMC, which is shown to have better overall performance, denoted DM-PMC [69,60].

For LAIS, we also test different number of proposals $N_{\text{LAIS}} \in \{10, 100, 200, 500\}$. We consider the *one-chain* application of LAIS (OC-LAIS), that requires to run one MCMC algorithm targeting $\bar{\pi}(\mathbf{x})$ to obtain the N_{LAIS} location parameters, hence it requires N_{LAIS} evaluations of the target. Then, at each iteration of LAIS, one sample is drawn from the mixture of proposals, hence we run the algorithm for $T_{\text{LAIS}} = 1000 - N_{\text{LAIS}}$ iterations for a fair comparison. For simplicity, we also consider Gaussian random-walk Metropolis to obtain the N_{LAIS} means.

Finally, we consider AMIS with several combinations of number of iterations T_{AMIS} and number of samples per iteration M . At each iteration, M samples are drawn from a single Gaussian proposal, hence the total number of evaluations is $E = MT_{\text{AMIS}}$. In this case, we test $E \in \{1000, 2000, 3000, 5000\}$, so the comparison is not fair (penalizing our approach) except for $E = 1000$.

Regarding our method, we use a mixture of $N_{\text{LAIS}} \in \{100, 200, 500\}$ proposal pdfs obtained by LAIS as $\bar{q}_{\text{par}}(\mathbf{x})$ as in Eq. (21) (we also use the means of these proposals as initial nodes). We vary N , and run our combined scheme for $T = \frac{E - N_{\text{LAIS}}}{N}$, keeping the number of target evaluations $E = 1000$. For PMC, LAIS and AMIS, as well as for the random walk proposal within the Metropolis algorithm, the covariance of the Gaussian proposals was set to $\xi^2\mathbf{I}_{10}$ and we test $\xi = 1, \dots, 6$. All the methods are compared through the mean absolute error (MAE) in estimating Z , and the results are averaged over 500 independent simulations.

The results are shown in Table 5, Table 6 and Table 7. We can see that NN-AIS+LAIS provides more robust results than only using LAIS. Namely, NN-AIS+LAIS obtains the same or a lower MAE than LAIS, depending on choice of the different parameters. Overall, the proposed scheme outperforms all the other benchmark AIS methods such as PMC, DM-PMC, LAIS and AMIS easily, even considering more target evaluations (penalizing our scheme) as shown in Table 7.

8.3. Inference in an astronomical model

In recent years, the problem of revealing objects that orbit other stars has acquired large attention in Astronomy. Different techniques have been proposed to discover exo-objects but, nowadays, the radial velocity technique is still the most used [70–72]. The model is highly non-linear and it is costly in terms of computation time (specially, for certain sets of parameters). The evaluation of the posterior involves numerically integrating a differential equation in time or an iterative procedure for solving a non-linear equation. Typically, the iteration is performed until a threshold is reached, or a certain number of iterations (e.g., typically 10^6 iterations), are performed. For the radial velocity model, this is needed for solving Eq. (36) described below. In the following, we describe an orbital model, which is equivalent for any N-body system observed from Earth, i.e. exoplanetary systems, binary stellar system, double pulsars, etc.

Table 6
MAE for Z with $E = 1000$ (best of each combination of N_{LAIS} and ξ are boldfaced).

Methods		$\xi = 1$	$\xi = 2$	$\xi = 3$	$\xi = 4$	$\xi = 5$	$\xi = 6$
NN-AIS+LAIS ($N_{\text{LAIS}} = 100$)	$N = 50$	0.9778	0.3886	0.1334	0.1487	0.1624	0.1968
	$N = 100$	0.9900	0.4152	0.1408	0.1519	0.1853	0.2502
	$N = 300$	0.9907	0.4817	0.1761	0.1466	0.1869	0.2427
NN-AIS+LAIS ($N_{\text{LAIS}} = 200$)	$N = 100$	0.7662	0.1607	0.1332	0.1179	0.1300	0.2000
	$N = 200$	0.8195	0.2176	0.1001	0.1250	0.1418	0.1854
	$N = 400$	0.8417	0.2954	0.1512	0.1218	0.1522	0.2060
NN-AIS+LAIS ($N_{\text{LAIS}} = 500$)	$N = 50$	0.2428	0.1801	0.1614	0.1313	0.1190	0.1642
	$N = 100$	0.2905	0.1406	0.1144	0.1046	0.1152	0.1851
	$N = 250$	0.4139	0.1270	0.1226	0.0989	0.1262	0.1783

Table 7
MAE for Z of AMIS with $E \in \{1000, 2000, 3000, 5000\}$. Note that the comparison is unfair (penalizing our approach) except for $E = 1000$.

Methods		$\xi = 1$	$\xi = 2$	$\xi = 3$	$\xi = 4$	$\xi = 5$	$\xi = 6$
AMIS $E = 1000$	$M = 10$	0.9998	0.9997	0.9997	0.9996	0.9996	0.9995
	$M = 100$	1.0000	1.0000	1.0000	0.9999	0.9997	0.9990
	$M = 200$	1.0000	1.0000	1.0000	1.0000	0.9998	0.9994
	$M = 500$	1.0000	1.0000	1.0000	1.0000	0.9998	0.9989
AMIS $E = 2000$	$M = 10$	0.9155	0.9117	0.8981	0.8987	0.8891	0.8878
	$M = 100$	0.9998	0.9986	0.9934	0.9784	0.9559	0.9072
	$M = 200$	1.0000	1.0000	0.9998	0.9981	0.9888	0.9712
	$M = 500$	1.0000	1.0000	1.0000	0.9998	0.9984	0.9953
AMIS $E = 3000$	$M = 10$	0.3293	0.3402	0.3051	0.3381	0.3540	0.3443
	$M = 100$	0.9725	0.9040	0.7963	0.6384	0.4964	0.3816
	$M = 200$	0.9998	0.9977	0.9884	0.9527	0.8308	0.7119
	$M = 500$	1.0000	1.0000	0.9998	0.9988	0.9859	0.9566
AMIS $E = 5000$	$M = 10$	0.0766	0.0768	0.0695	0.0722	0.0699	0.0725
	$M = 100$	0.1626	0.1176	0.0957	0.0810	0.0737	0.0656
	$M = 200$	0.8771	0.6040	0.2824	0.1473	0.1163	0.0899
	$M = 500$	1.0000	0.9982	0.9904	0.9449	0.7944	0.4532

Table 8
Description of parameters in Eq. (33).

Parameter	Description	Units
For each planet		
ζ_i	amplitude of the curve	m s^{-1}
$u_{i,k}$	true anomaly	rad
ω_i	longitude of periastron	rad
e_i	orbit's eccentricity	...
P_i	orbital period	s
τ_i	time of periastron passage	s
Below: not depending on the number of objects/satellite		
V_0	mean radial velocity	m s^{-1}

8.3.1. Likelihood function and prior densities

When analysing radial velocity data of an exoplanetary system, it is commonly accepted that the *wobbling* of the star around the center of mass is caused by the sum of the gravitational force of each planet independently and that they do not interact with each other. Each planet follows a Keplerian orbit and the radial velocity of the host star is given by

$$y_k = V_0 + \sum_{i=1}^S \zeta_i [\cos(u_{i,k} + \omega_i) + e_i \cos(\omega_i)] + \xi_k, \quad (33)$$

with $k = 1, \dots, K$. The number of objects in the system is S , that is consider known in this experiment (for the sake of simplicity). Note that the iteration index $i = 1, \dots, S$ denotes the i -th object/planet. Both $y_k, u_{i,k}$ depend on time t , and ξ_k is a Gaussian noise perturbation with variance σ_e^2 . For simplicity, we consider this value known, $\sigma_e^2 = 1$. The meaning of each parameter in Eq. (33) is given in Table 8. The likelihood function is defined by (33) and some indicator variables described below. The angle $u_{i,k}$ is the true anomaly of the planet i and it can be determined from

$$\frac{du_{i,k}}{dt} = \frac{2\pi}{P_i} \frac{(1 + e_i \cos u_{i,k})^2}{(1 - e_i)^{\frac{3}{2}}} \quad (34)$$

This equation has analytical solution. As a result, the true anomaly $u_{i,k}$ can be determined from the mean anomaly $M_{i,k}$. However, the analytical solution contains a non linear term that needs to be determined by iterating. First, we define the mean anomaly $M_{i,k}$ as

$$M_{i,k} = \frac{2\pi}{P_i} (t - \tau_i), \quad (35)$$

where τ_i is the time of periastron passage of the planet i and P_i is the period of the orbit (see Table 8). Then, through the Kepler's equation,

$$M_{i,k} = E_{i,k} - e_i \sin E_{i,k}, \quad (36)$$

we have to obtain $E_{i,k}$, which is the eccentric anomaly. Equation (36) has no analytic solution and it must be solved by an iterative procedure. A Newton-Raphson method is typically used to find the roots of this equation [73]. For certain sets of parameters this iterative procedure can be particularly slow.

Finally, we can also obtain $u_{i,k}$ from

$$\tan \frac{u_{i,k}}{2} = \sqrt{\frac{1 + e_i}{1 - e_i}} \tan \frac{E_{i,k}}{2}. \quad (37)$$

Hence, the vector of variables to infer, \mathbf{x} , is

$$\mathbf{x} = [V_0, \zeta_1, \omega_{1,t}, e_1, P_1, \tau_1, \dots, \zeta_S, \omega_S, e_S, P_S, \tau_S]. \quad (38)$$

For a single object (e.g., a planet or a natural satellite), the dimension of \mathbf{x} is $d_x = 5 + 1 = 6$, with two objects the dimension of \mathbf{x} is $d_x = 11$ etc. Generally, we have $d_x = 1 + 5S$. Note that the observation model in Eq. (33) induces the likelihood function $p(\mathbf{y}|\mathbf{x})$, where $\mathbf{y} = [y_1, \dots, y_K]$.

Table 9
Relative Mean Square Errors (MSE) and normalized computational time.

Methods	NN-AIS+U	IS	IS	IS	IS
RMSE	5.755	9.439	7.943	6.524	5.431
normalized time	1.53	1	1.91	3.20	4.17
posterior evaluations (E)	10^6	10^6	$2 \cdot 10^6$	$3 \cdot 10^6$	$4 \cdot 10^6$

Priors As prior densities we consider uniform pdfs in the following intervals: $V_0 \in [-20, 20]$, $\zeta_i \in [0, 50]$, $e_i \in [0, 1]$, $P_i \in [0, 365]$, $\omega_{i,k} \in [0, 2\pi]$, $\tau_i \in [0, P_i]$ (i.e., the prior is zero outside these intervals), for all $i = 1, \dots, S$. This means that the likelihood function is zero when the particles fall out of these intervals. Note that the interval of τ_i is conditioned to the value P_i . This parameter is the time of periastron passage, i.e. the time passed since the object passed the closest point in its orbit. It has the same units of P_i and can take values from 0 to P_i .

8.3.2. Experiment setting and results

We generate a set of data $\{y_k\}_{k=1}^K$ with $K = 50$, and $S = 2$ objects (so that $d_x = 11$), according to the observation model above. We set $V = 2$, $\zeta_1 = 25$, $\omega_1 = 0.61$, $e_1 = 0.1$, $P_1 = 15$, $\tau_1 = 3$ (for the first object) and $\zeta_2 = 5$, $\omega_2 = 0.17$, $e_2 = 0.3$, $P_2 = 115$, $\tau_2 = 25$ (for the second object). We compare a standard IS scheme using the prior as proposal and the NN-AIS+U scheme (using again the prior as uniform proposal component) using the parsimonious scheme with acceptance function A3 in Eq. (27). In NN-AIS+U, we consider $N = 10000$, $T = 100$ and $L = 10^6$. The total number of evaluations of the posterior is then $NT = 10^6$ for NN-AIS+U. For the standard IS scheme, we consider different number of samples $\{10^6, 2 \cdot 10^6, 3 \cdot 10^6, 4 \cdot 10^6\}$. We compute the Relative MSE (RMSE) in estimation of the 11 parameters in \mathbf{x} , averaged over all the components. The results are also averaged over 200 independent runs. Table 9 provides the RMSE and the computational time, normalized with respect to the time spent by the standard IS scheme with 10^6 samples. We can observe that, in order to obtain the same performance of NN-AIS+U in terms of RMSE, the IS schemes require much more computational time than NN-AIS+U. Therefore, this is an example with a real-world model where the inequality (22) is fulfilled.

8.4. Retrieval of biophysical parameters inverting an RTM model

In this experiment, we apply NN-AIS to retrieve biophysical parameters of a sequence of problems involving the radiative transfer PROSAIL model. The purpose is to show the ability of NN-AIS to share information from related inverse problems easily. The combined PROSPECT leaf optical properties model and SAIL canopy bidirectional reflectance model, also referred to as PROSAIL, have been used for almost two decades to study plant canopy spectral and directional reflectance in the solar domain [74]. PROSAIL has also been used to develop new methods for retrieval of vegetation biophysical properties. It links the spectral variation of canopy reflectance, which is mainly related to leaf biochemical contents, with its directional variation, which is primarily related to canopy architecture and soil/vegetation contrast. This link is key to simultaneous estimation of canopy biophysical/structural variables for applications in agriculture, plant physiology, and ecology at different scales. PROSAIL has become one of the most popular radiative transfer tools due to its ease of use, general robustness, and consistent validation by lab/field/space experiments over the years.

Inversion of PROSAIL The context is Bayesian inversion of an observation model $\mathbf{h}(\mathbf{x})$.³ In our setting, the observation model is

Table 10
Description of parameters in Eq. (39).

Parameter	Description	Units
S_{st}	structure coefficient	-
C_{hl}	chlorophyll content	$\mu\text{g cm}^{-2}$
C_{ar}	carotenoid content	$\mu\text{g cm}^{-2}$
C_{br}	brown pigment content	-
C_w	water content	cm
C_m	dry matter content	g cm^{-2}

Table 11
Characteristics of the simulation used in the PROSAIL model.

Canopy level	LAI	ALA	Hotspot	θ_s	θ_v	$\Delta\Theta$
5	30	0.01	30	10	90	

PROSAIL, which models reflectance in terms of leaf optical properties and canopy level characteristics. We choose only leaf optical properties as the set parameters of interest

$$\mathbf{x} = [S_{st}, C_{hl}, C_{ar}, C_{br}, C_w, C_m] \in \mathbb{R}^6, \tag{39}$$

described in Table 10. In Table 11, we show the fixed values of canopy level characteristics, which are determined by the leaf area index (LAI), the average leaf angle inclination (ALA), the hot-spot parameter (Hotspot), and the parameters of system geometry described by the solar zenith angle (θ_s), view zenith angle (θ_v), and the relative azimuth angle between both angles ($\Delta\Theta$). The observation model is $\mathbf{y} = \mathbf{h}(\mathbf{x}) + \mathbf{v}$, where $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_{d_y})$ with $\sigma = 1$. The observed data, denoted $\mathbf{y} \in \mathbb{R}^{d_y}$ with $d_y = 2101$, corresponds to the detected spectra. We generated synthetic spectra and the goal is to infer \mathbf{x} studying the corresponding posterior distribution. The Gaussian noise $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I})$ jointly with PROSAIL, $\mathbf{h}(\mathbf{x})$, induces the following likelihood function

$$\ell(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{y}|\mathbf{h}(\mathbf{x}), \sigma^2 \mathbf{I}). \tag{40}$$

We set the prior $g(\mathbf{x})$ as a product of indicator variables $S_{st} \in [1, 3]$, $C_{hl} \in [0, 100]$, $C_{ar} \in [0, 25]$, $C_{br} \in [0, 1]$, $C_w \in [0, 0.05]$ and $C_b \in [0, 0.02]$, i.e., the prior is zero outside these intervals.⁴ The complete posterior is then $p(\mathbf{x}|\mathbf{y}) = \frac{1}{Z} \ell(\mathbf{y}|\mathbf{x})g(\mathbf{x})$. It is important to remark that PROSAIL is an *highly non-linear* model and its inversion is a very complicated problem, as shown the remote sensing literature [75,67].

Sequential inversion for image recovery In remote sensing, the goal is usually to recover of an image formed by R pixels. A set of physical parameters \mathbf{x}_r is associated to the r -th pixel. Hence, the corresponding vector of observations \mathbf{y}_r is also associate to each pixel. We have then a collection of inverse problems, where we desire to retrieve \mathbf{x}_r given \mathbf{y}_r , one for each pixel. Mathematically, let consider R measurements, $\{\mathbf{y}_r\}_{r=1}^R$, associated each to a different inverse problem, under the PROSAIL model, i.e., a mapping $\mathbf{h}(\mathbf{x}) : \mathbb{R}^{d_x} \rightarrow \mathbb{R}^{d_y}$,

$$\mathbf{y}_r = \mathbf{h}(\mathbf{x}_r) + \mathbf{v}_r, \quad r = 1, \dots, R. \tag{41}$$

We assume $\mathbf{v}_r \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_{d_y})$, for all $r = 1, \dots, R$, with $d_y = 2101$, and $\sigma = 1$, and thus we have a R posterior distributions $p_r(\mathbf{x}_r|\mathbf{y}_r)$ for $r = 1, \dots, R$ (we recall that $\mathbf{x}_r \in \mathbb{R}^{d_x}$, with $d_x = 6$). We solve them sequentially while reusing information. Some examples of data \mathbf{y}_r and model values are given in Fig. 12.

³ The MATLAB code of PROSAIL is available in <http://teledetection.ipgp.jussieu.fr/prosail/>.

⁴ We have employed the ranges suggested <http://opticleaf.ipgp.fr/index.php?page=prospect>.

Table 12
Relative Mean Absolute Errors (RMAE) for each component (averaged over all spectra).

Components	1	2	3	4	5	6	Mean
Stand. IS	0.7556	0.4397	2.9431	0.6247	0.2096	0.2782	2.8516
Sequential NN-AIS	0.2045	0.2245	0.8891	0.1985	0.1425	0.1320	1.0715

Table 13
Mean Absolute Errors (RMAE) for each component (averaged over all spectra).

Components	1	2	3	4	5	6	Mean
Stand. IS	0.9760	6.1754	9.8204	0.1348	0.0016	0.0012	0.8752
Sequential NN-AIS	0.2641	3.1535	2.9667	0.0428	0.0011	0.0006	0.2985

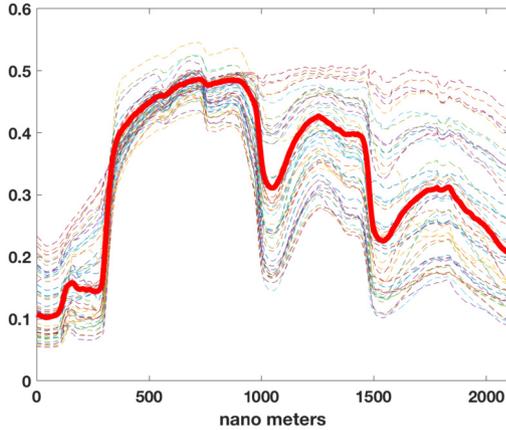


Fig. 12. An example of vector of data \mathbf{y}_r (hyperspectral reflectances, shown with solid line) and the model values corresponding to 50 different samples, $\mathbf{f}^{(i)} = \mathbf{f}(\mathbf{x}_r^{(i)})$ (dashed lines). Each component of the vector \mathbf{y}_r corresponds to a different wavelength (nm).

Experiment In a real data settings, physical and geographical patterns are associated to the parameters \mathbf{x}_r in the image. In order to check the performance of each algorithm, we consider synthetic data. Thus, in this experiment, we have also generated synthetic patterns in order to simulate a real scenario. In particular, we produce six patterns (recall $\mathbf{x} \in \mathbb{R}^6$) that represent handwritten digits (see Fig. 13). Hence, in this setting, we have $R = 784$ different observation vectors \mathbf{y}_r , $r = 1, \dots, R$, for which we want to estimate the vectors of true values \mathbf{x}_r , $r = 1, \dots, R$. Each observation corresponds to a single pixel of a 28×28 image. We also compute the maximum a-posteriori (MAP) of $p_r(\mathbf{x}_r | \mathbf{y}_r)$, $\mathbf{x}_{r, \text{MAP}}$, as estimate of \mathbf{x}_r .

Methods We use the NN-AIS scheme to estimate $\mathbf{x}_{r, \text{MAP}}$ for $r = 1, \dots, 784$, and compare it against IS using the prior as proposal density, in terms of relative squared error and by looking at the recovered images. As parameters of our scheme we chose $N_{\text{init}} = 1000$, $T = 20$, $N = 250$ and $L = 10^5$. The N_{init} initial points were taken at random in the domain except for 11 points that were placed in the vertices of the domain. Our scheme allows for sharing information from problem to the next one, so we also use the $\hat{\mathbf{x}}_{s, \text{MAP}}$ for $s = 1, \dots, r - 1$ as initial nodes when estimating $\mathbf{x}_{r, \text{MAP}}$. Note that this is completely fair since the model has been already evaluated at those points. The comparison is fair in terms of model evaluations, with a total of $E = 6000$ for each $r = 1, \dots, 784$.

Results The results are shown in Figs. 14 and 15. It can be seen that both standard IS and NN-AIS are able to correctly recover components 2, 4, 5 and 6 of \mathbf{x}_r ($r = 1, \dots, 784$), i.e., the images of “2”, “4”, “5” and “6” in both Fig. 14 and Fig. 15 look very close to the true ones (Figs. 13(b), (d), (e) and (f) respectively). The images recovered by NN-AIS have lower noise though. The components 1 and 3 of the \mathbf{x}_r 's are completely lost with standard IS (see Fig. 14), whereas NN-AIS is able at least to achieve to recover the bound-

Table 14
Absolute and relative error (averaged over all the pixels) in the transformed domain (“reconstruction of the spectra”).

	Absolute	Relative
Stand. IS	66.4395	0.0802
Sequential NN-AIS	11.0844	0.0198

aries of the corresponding patterns. Indeed, NN-AIS obtains a much lower error in estimation, as it is shown in Table 12 and Table 13. The difficulty in recovering the components 1 (i.e., S_{st}) and 3 (i.e., C_{ar}) deserves further studies. This issue could be related to some relevant features of PROSAIL (e.g., the average partial derivatives with respect to these two components). We leave the study of these specific issues for future work. In Table 14, we also show the averaged error in the spectra produced by both methods as compared to the true observations.

9. Conclusions and future lines

In this work, we introduced a novel framework of adaptive importance sampling algorithms. The key idea is the use of a non-parametric proposal density built by a regression procedure (the emulator), that mimics the true shape of posterior pdf. Hence, the proposal pdf represents also a surrogate model, that is in turn adapted through the iterations by adding new support points. The regression (e.g., obtained by nearest neighbors and Gaussian processes) can be applied directly on the posterior domain or, alternatively, in just one piece of the likelihood, such as an arbitrary physical model. Drawing from the emulator is possible by a deep architecture of two nested IS layers. More sophisticated deep structures, employing a chain of emulators, have been described.

RADIS is an extremely efficient importance sampling scheme since the emulator (used as proposal pdf) becomes closer and closer to the true posterior, as new nodes are incorporated. As a consequence, RADIS asymptotically converges to an exact sampler under mild conditions. Several numerical experiments and theoretical supports confirm these statements. Robust accelerating versions of RADIS have been also presented, as well as combinations with other benchmark AIS algorithms. Cheap constructions of the emulator have been also discussed and tested. The use of RADIS within a sequential Monte Carlo scheme will be considered in future works. Furthermore, as future research lines, we also plan to analyze in depth the PROSAIL inversion problem, approximating the partial derivatives with respect some specific parameters by RADIS. Moreover, we also plan to consider the adaptation of the auxiliary proposal $\tilde{q}_{\text{aux}}(\mathbf{x})$, adding also additional layers in the proposed deep architecture.

CRediT authorship contribution statement

The authors contributed in equal way to all the manuscript.

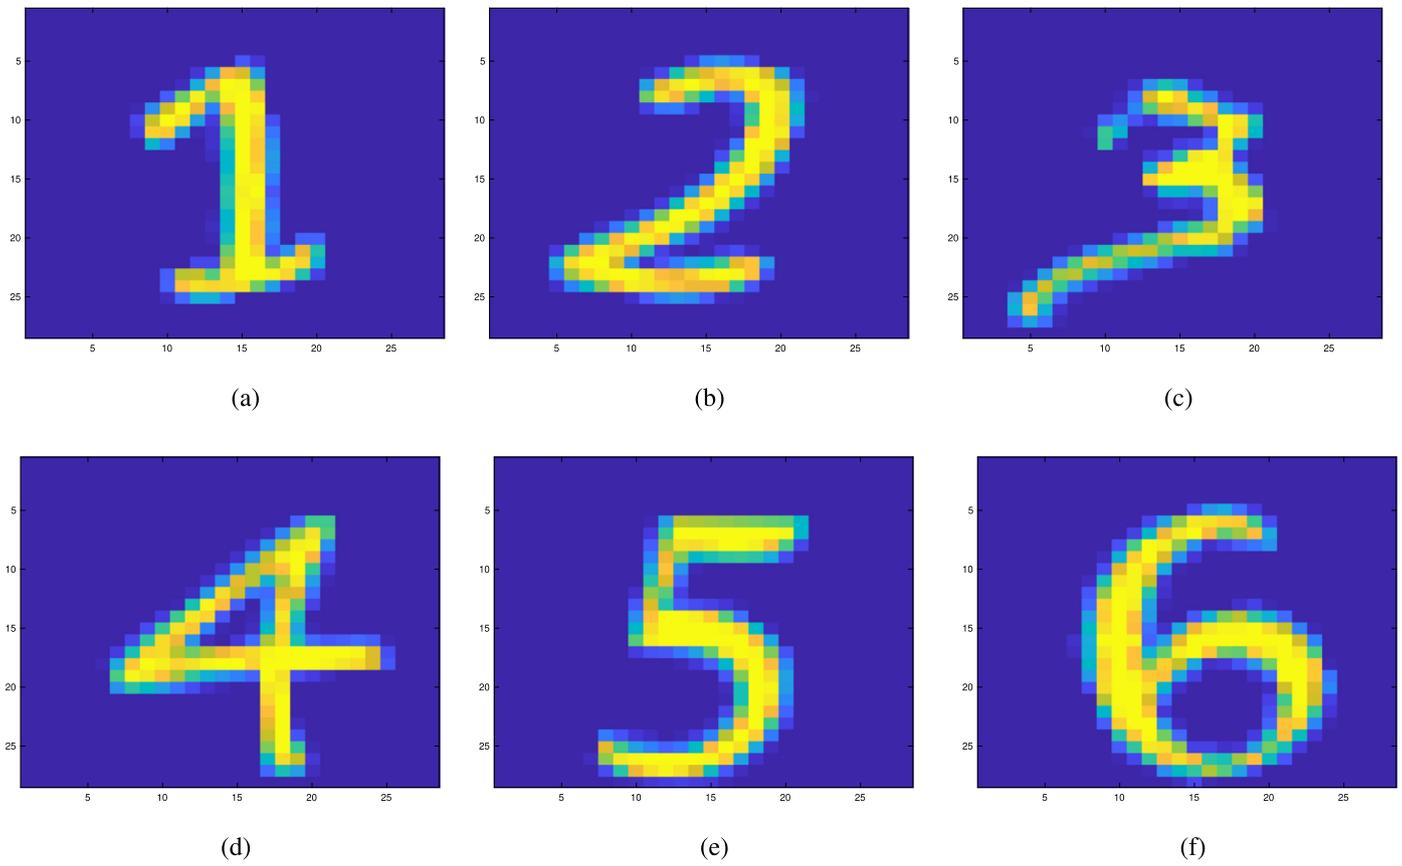


Fig. 13. Patterns of the true parameter values (scaled according to range of each parameter), i.e., the ground-truths.

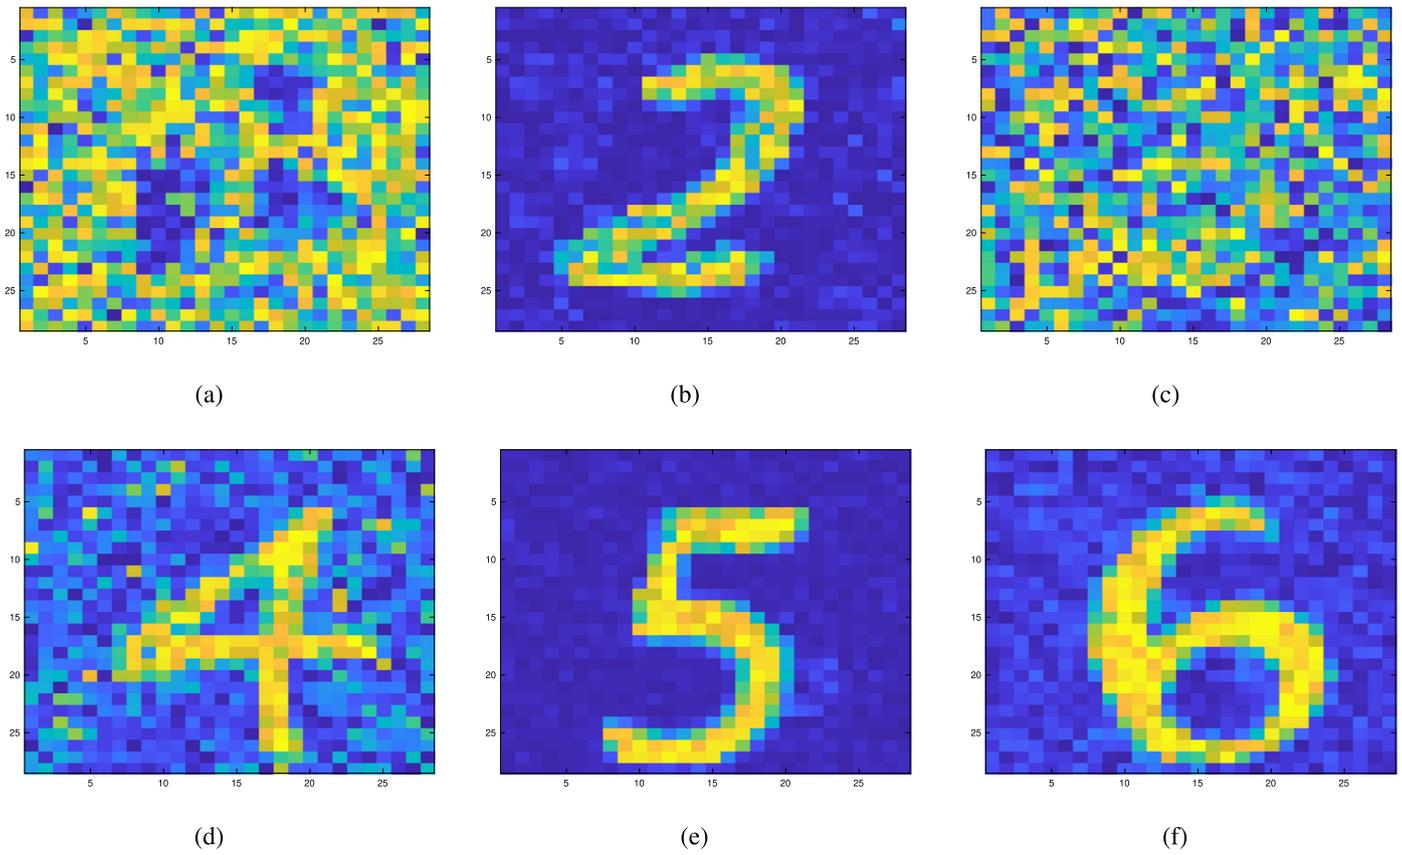


Fig. 14. Recovered by standard IS. We can observe the difficulty in the retrieval of the first and third parameter.

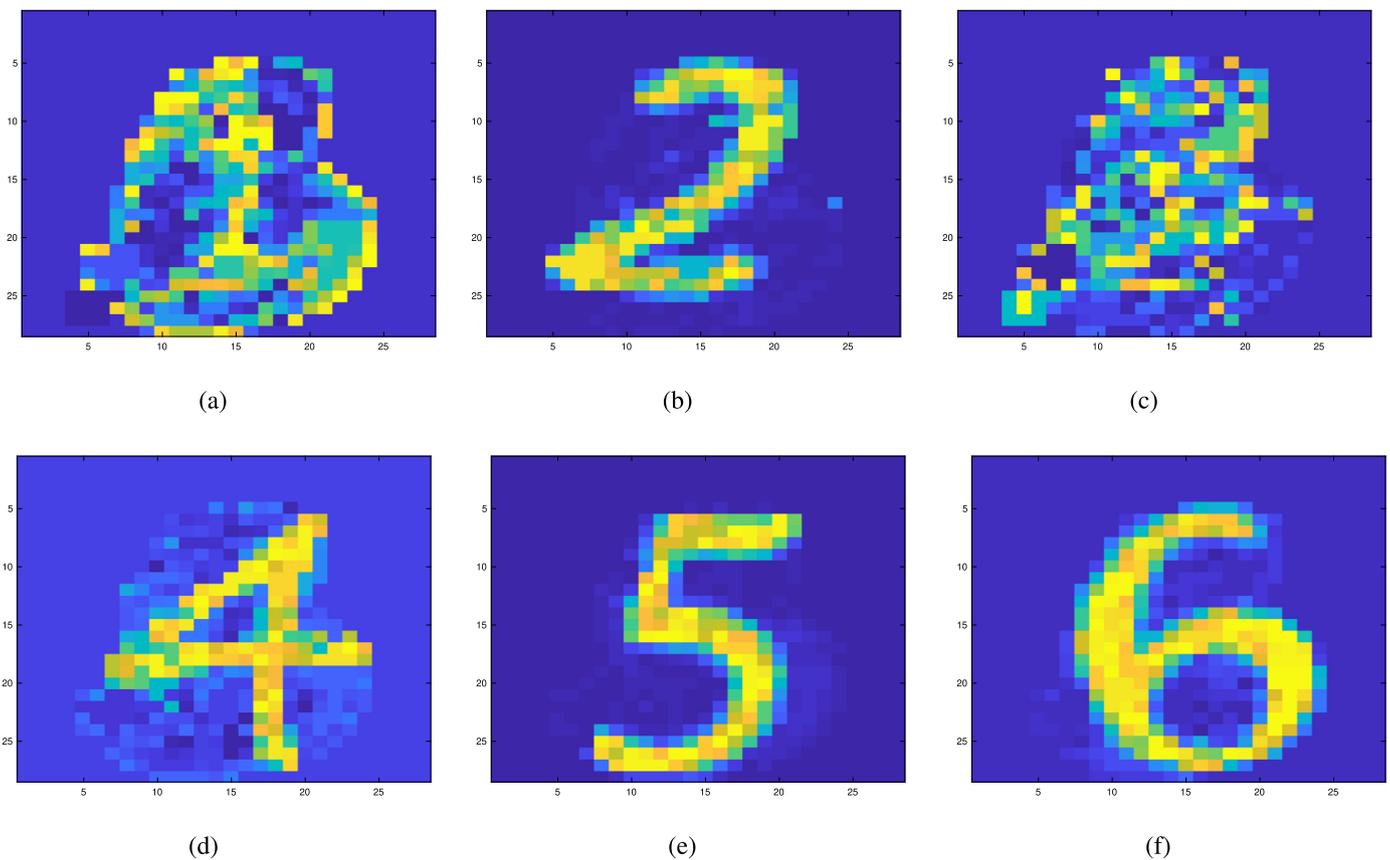


Fig. 15. Recovered by NN-AIS. We can observe the retrieval of first and third parameter is not completely successful.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work has been supported by Spanish government via grant FPU19/00815, by Agencia Estatal de Investigación AEI (project SP-GRAPH, ref. num. PID2019-105032GB-I00), by the Found action by the Community of Madrid in the framework of the Multiannual Agreement with the Rey Juan Carlos University in line of action 1, “Encouragement of Young Phd students investigation”, Project Ref. F661 Acronym Mapping-UCI, and by the European Research Council (ERC) under the ERC Consolidator Grant 2014 project SEDAL (647423).

Appendix A. Theoretical support

In this section, we discuss several theoretical aspects of RADIS. First, we address the error in the approximate sampling and evaluation of the interpolating proposal. Then, we show that the adaptive construction of the proposal decreases the distance with respect to the true target as the number of nodes J_t grows. Finally, we show that this also minimizes the variance of the IS weights.

A.1. Sampling Importance Resampling (SIR)

Let $\hat{\pi}_t(\mathbf{x})$ the unnormalized interpolating proposal from which we aim to sample. Its normalizing constant c_t is not important in this first part. The SIR method allows to sample from the density

$\hat{\pi}_t$ by resampling a sample drawn from another auxiliary (importance) density [54][76, Chapter 24]. This method is also referred as the weighted bootstrap in [55, Sect. 3.2]. The SIR algorithm is as follows:

1. Draw $\{\mathbf{x}_1, \dots, \mathbf{x}_L\}$ i.i.d. from $\bar{q}_{aux}(\mathbf{x})$, that is a density with fatter tails than $\hat{\pi}_t(\mathbf{x})$.
2. Calculate the importance weights for each \mathbf{x}_i

$$\gamma_i = \gamma(\mathbf{x}_i) = \frac{\hat{\pi}_t(\mathbf{x}_i)}{\bar{q}_{aux}(\mathbf{x}_i)}.$$

3. Resample N ($N \leq L$) values $\{\mathbf{x}_1^*, \dots, \mathbf{x}_N^*\}$ from $\{\mathbf{x}_1, \dots, \mathbf{x}_L\}$ with probabilities proportional to γ_i assigned to \mathbf{x}_i .

If $L \rightarrow \infty$, or more precisely $\frac{L}{N} \rightarrow \infty$, then the set $\{\mathbf{x}_1^*, \dots, \mathbf{x}_N^*\}$ is asymptotically distributed as $\hat{\pi}_t(\mathbf{x})$. Thus, the choice of L and N is important for two factors: (i) to reduce the dependence of the \mathbf{x}_i^* s, and (ii) to have the distribution of \mathbf{x}_i^* as close to $\hat{\pi}_t$ as possible. The relative magnitude between N and L controls this dependence, while only the magnitude of L affects how close the distribution of the resampled particle is to the density $\hat{\pi}_t$.

Bias and correlation in SIR Under mild conditions, as $\frac{L}{N} \rightarrow \infty$, the density of resampled particle converges to $\hat{\pi}_t(\mathbf{x})$. For more details see [56, Sect. 6.2.4], [55, Sect. 3.2] and [57, page 6]. As SIR is an approximate sampling algorithm, it has some bias.⁵ If the first and second moment of the IS weight $\gamma(\mathbf{x}) = \frac{\hat{\pi}_t(\mathbf{x}_i)}{\bar{q}_{aux}(\mathbf{x})}$ exists, it can be shown that this bias vanishes at $\mathcal{O}(L^{-1})$ rate [76, Chapter 24]. In

⁵ Measured as the difference in the probability of some set between the target pdf and the “equivalent” pdf.

[55, Sect. 3.2], they show the convergence of the cdf of the resampled particle as $L \rightarrow \infty$ in the univariate case.

Resampling N times from a unique pool of L samples from $\bar{q}_{\text{aux}}(\mathbf{x})$ introduces correlation in the resampled sample. However, when $N \ll L$, this correlation is negligible. Some heuristics suggest $\frac{L}{N} = 20$ [54], or $\frac{L}{N} \geq 10$ [55]. For more details in the relation of the values of L and N see [76, Sect. 24.3]. In [58] (see Figure 5 and Appendix A therein), it is shown the “equivalent” density of a resampled particle for a fixed value of L , which converges to the target pdf as L diverges. Furthermore, for computing the denominator of the outer weights we need the normalizing constant of $\hat{\pi}_t$ (see Eq. (11)). In this sense, the inner IS also provides with an approximation by using the L samples from $\bar{q}_{\text{aux}}(\mathbf{x})$,

$$\hat{c}_t = \frac{1}{L} \sum_{\ell=1}^L \frac{\hat{\pi}_t(\mathbf{z}_\ell)}{\bar{q}_{\text{aux}}(\mathbf{z}_\ell)}. \quad (\text{A.1})$$

This estimate converges as $L \rightarrow \infty$ [3].

A.2. Variance of the IS weights

Let $w(\mathbf{x}) = \frac{\pi(\mathbf{x})}{\bar{q}(\mathbf{x})}$ be the weight function evaluated at samples $\mathbf{x} \sim \bar{q}(\mathbf{x})$. First of all, note that $E[w(\mathbf{x})] = Z$. Below, we show that the variance of $w(\mathbf{x})$ is proportional to the Pearson divergence between the posterior $\bar{\pi}$ and proposal q , i.e.,

$$\text{var}[w(\mathbf{x})] = \int_{\mathcal{X}} (w(\mathbf{x}) - Z)^2 \bar{q}(\mathbf{x}) d\mathbf{x} \quad (\text{A.2})$$

$$= \int_{\mathcal{X}} \left(\frac{\pi(\mathbf{x}) - Z\bar{q}(\mathbf{x})}{\bar{q}(\mathbf{x})} \right)^2 \bar{q}(\mathbf{x}) d\mathbf{x} \quad (\text{A.3})$$

$$= Z^2 \int_{\mathcal{X}} \frac{(\bar{\pi}(\mathbf{x}) - \bar{q}(\mathbf{x}))^2}{\bar{q}(\mathbf{x})} d\mathbf{x} = Z^2 \chi^2(\bar{\pi} \| q), \quad (\text{A.4})$$

where $\chi^2(\bar{\pi} \| q) = \int_{\mathcal{X}} \frac{(\bar{\pi}(\mathbf{x}) - \bar{q}(\mathbf{x}))^2}{\bar{q}(\mathbf{x})} d\mathbf{x}$, is the Pearson divergence and we have used $\bar{\pi}(\mathbf{x}) = \frac{1}{Z} \pi(\mathbf{x})$. Hence, if we construct a proposal such $\chi^2(\bar{\pi} \| q) \rightarrow 0$, we would obtain $\text{var}[\hat{Z}] = 0$. Moreover, the mean square error (MSE) of \hat{T} can also be shown to be bounded by this divergence (see e.g. [53])

$$\mathbb{E}[|I - \hat{T}|^2] \leq \frac{C_f (\chi^2(\bar{\pi} \| \bar{q}) + 1)}{N}. \quad (\text{A.5})$$

Thus, it is beneficial to reduce the $\chi^2(\bar{\pi} \| \bar{q})$ in order to obtain accurate IS estimators.

A.3. Pearson divergence and L_p distances

Now, we aim to show that $\chi^2(\bar{\pi} \| \bar{q})$ can be bounded in terms of the L_2 and L_∞ distances, between $\bar{\pi}(\mathbf{x})$ and $\bar{q}(\mathbf{x})$. Using Holder's inequality and the fact that pdfs are always positive, we can write

$$\begin{aligned} \chi^2(\bar{\pi} \| \bar{q}) &= \int_{\mathcal{X}} |\bar{\pi}(\mathbf{x}) - \bar{q}(\mathbf{x})| \frac{|\bar{\pi}(\mathbf{x}) - \bar{q}(\mathbf{x})|}{|\bar{q}(\mathbf{x})|} d\mathbf{x} \\ &= \left\| (\bar{\pi} - q) \left(\frac{\bar{\pi} - \bar{q}}{\bar{q}} \right) \right\|_{L_1} \\ &\leq \|\bar{\pi} - \bar{q}\|_{L_2} \left\| \frac{\bar{\pi} - \bar{q}}{\bar{q}} \right\|_{L_2}. \end{aligned} \quad (\text{A.6})$$

The L_2 distance can be easily shown to be bounded by L_∞ distance (considering a bounded domain \mathcal{X}), i.e.,

$$\begin{aligned} \|\bar{\pi} - \bar{q}\|_{L_2} &= \left(\int_{\mathcal{X}} |\bar{\pi}(\mathbf{x}) - \bar{q}(\mathbf{x})|^2 d\mathbf{x} \right)^{\frac{1}{2}} \\ &\leq \left(|\mathcal{X}| \max |\bar{\pi}(\mathbf{x}) - \bar{q}(\mathbf{x})|^2 \right)^{\frac{1}{2}} \\ &= |\mathcal{X}|^{\frac{1}{2}} \|\bar{\pi} - \bar{q}\|_{L_\infty}. \end{aligned} \quad (\text{A.7})$$

Similarly, we have

$$\left\| \frac{\bar{\pi} - \bar{q}}{\bar{q}} \right\|_{L_2} \leq |\mathcal{X}|^{\frac{1}{2}} \left\| \frac{\bar{\pi} - \bar{q}}{\bar{q}} \right\|_{L_\infty}. \quad (\text{A.8})$$

Thus, we can obtain the following result regarding the L_∞ distance,

$$\chi^2(\bar{\pi} \| \bar{q}) \leq |\mathcal{X}| \left\| \frac{\bar{\pi} - \bar{q}}{\bar{q}} \right\|_{L_\infty} \|\bar{\pi} - \bar{q}\|_{L_\infty}. \quad (\text{A.9})$$

Since we choose \bar{q} in order to have fatter tails than $\bar{\pi}$ and since $\bar{q}, \bar{\pi}$ are bounded, then the factor $\left\| \frac{\bar{\pi} - \bar{q}}{\bar{q}} \right\|_{L_\infty}$ in (A.9) vanishes to zero if $\|\bar{\pi} - \bar{q}\|_{L_\infty} \rightarrow 0$. Therefore, if $\|\bar{\pi} - \bar{q}\|_{L_\infty} \rightarrow 0$, we have $\chi^2(\bar{\pi} \| \bar{q}) \rightarrow 0$. Due to (A.7)-(A.8), this result is also valid for the L_2 distance. In this work, we consider $\bar{q} = \bar{q}_t = \frac{1}{c_t} \hat{\pi}_t$ such $\|\bar{\pi} - \hat{\pi}_t\|_{L_\infty} \rightarrow 0$ as $t \rightarrow \infty$ (see section below), and thus $\|\bar{\pi} - \frac{1}{c_t} \hat{\pi}_t\|_{L_\infty} \rightarrow 0$, that implies $\chi^2(\bar{\pi} \| \frac{1}{c_t} \hat{\pi}_t) \rightarrow 0$.

A.4. Convergence of the emulator to target function

For simplicity, let us focus on the interpolation case and a bounded \mathcal{X} . Here, we show that the interpolating constructions of Sect. 4.2, jointly with the adaptation process, lead to a proposal $\frac{1}{c_t} \hat{\pi}_t(\mathbf{x})$ that converges to $\bar{\pi}(\mathbf{x})$. Since \hat{c}_t is an unbiased estimation of the area below $\hat{\pi}_t(\mathbf{x})$, we focus on the convergence of $\hat{\pi}_t(\mathbf{x})$ to the unnormalized posterior $\pi(\mathbf{x})$. As $\hat{\pi}_t(\mathbf{x}) \rightarrow \pi(\mathbf{x})$, then $\hat{c}_t \rightarrow Z$. In Sect. 5, we have introduced an extra parametric density $\bar{q}_{\text{par}}(\mathbf{x})$ to also ensure that new points can be added in any region of the domain \mathcal{X} during the adaptation. We show below that, when using the NN or GP constructions, the approximation error of $\hat{\pi}_t$ depends on a quantity called *fill distance*,

$$r_t = \max_{\mathbf{x}} \min_{i=1, \dots, J_t} \|\mathbf{x} - \mathbf{x}_i\|_2, \quad (\text{A.10})$$

which measures the filling of the space. In other words, the greater the fill distance, the less covered the space is. For both constructions, decreasing the fill distance ensures that $\hat{\pi}_t(\mathbf{x})$ converges in L_∞ norm to $\pi(\mathbf{x})$. Using a $\bar{q}_{\text{par}}(\mathbf{x})$ that is not negative in \mathcal{X} ensures every region will be covered eventually, i.e., $r_t \rightarrow 0$ as $t \rightarrow \infty$.

NN construction If π is Lipschitz continuous, we have that

$$\|\pi - \hat{\pi}_t\|_\infty \leq L_0 r_t, \quad (\text{A.11})$$

where L_0 is the Lipschitz constant and r_t denotes the fill distance [21, App. D.4]. Equivalently, we have [77]

$$\|\pi - \hat{\pi}_t\|_\infty \leq L_0 \max_{i=1, \dots, J_t} \text{diam}(\mathcal{R}_i), \quad (\text{A.12})$$

that is, the approximation error is bounded by the biggest Voronoi cell. Covering the space (not necessarily with uniform points) ensure that $\max_i \text{diam}(\mathcal{R}_i) \rightarrow 0$ [78] (equivalently $r_t \rightarrow 0$), and thus $\hat{\pi}_t \rightarrow \pi$ as $t \rightarrow \infty$.

GP construction First, we recall a result valid when the GP regression is applied on π , not a transformation. It can be shown that the approximation error $\|\pi - \hat{\pi}_t\|_\infty$ is bounded in terms of the fill distance (e.g. see [21, Sect. 7] and references therein)

$$\|\pi - \hat{\pi}_t\|_\infty = \mathcal{O}(\lambda(r_t)). \quad (\text{A.13})$$

The speed of convergence, i.e., the functional form of $\lambda(r_t)$, depends on the choice of kernel (e.g. under some circumstances and with Gaussian kernel, $\lambda(r_t)$ decays exponentially when $r_t \rightarrow 0$).

In case we do not approximate $\pi(\mathbf{x})$ directly, but we build an emulator of $\log \pi(\mathbf{x})$ or just on the physical model $\mathbf{h}(\mathbf{x})$, it is also possible to show the convergence of the posterior approximation. See, for instance, the error bounds in [79, Theorem 4.2].

Appendix B. A special interesting case for NN-AIS

Here, we focus on NN-AIS. We consider a bounded \mathcal{X} and building $\hat{\pi}_t$ with a nearest neighbor (NN) approach. In Sect. 4.2, we show that the NN emulator at iteration t is given by

$$\hat{\pi}_t(\mathbf{x}) = \sum_{i=1}^{J_t} \pi(\mathbf{x}_i) \mathbb{I}_{\mathcal{R}_i}(\mathbf{x}) = \sum_{i=1}^{J_t} \pi(\mathbf{x}_i) |\mathcal{R}_i| \left[\frac{1}{|\mathcal{R}_i|} \mathbb{I}_{\mathcal{R}_i}(\mathbf{x}) \right], \quad (\text{B.1})$$

$$= \sum_{i=1}^{J_t} v_i p_i(\mathbf{x}), \quad (\text{B.2})$$

where $|\mathcal{R}_i|$ is the measure of i -th Voronoi region (see Eq. (18) for the definition of \mathcal{R}_i), $v_i = \pi(\mathbf{x}_i) |\mathcal{R}_i|$, and $p_i(\mathbf{x}) = \frac{1}{|\mathcal{R}_i|} \mathbb{I}_{\mathcal{R}_i}(\mathbf{x})$ are uniform densities over \mathcal{R}_i . Hence, $\hat{\pi}_t(\mathbf{x})$ is a mixture of J_t uniform densities where the mixture weight is proportional to v_i . The normalizing constant of $\hat{\pi}_t(\mathbf{x})$ is given by

$$c_t = \sum_{i=1}^{J_t} v_i = \sum_{i=1}^{J_t} \pi(\mathbf{x}_i) |\mathcal{R}_i|, \quad (\text{B.3})$$

so that the normalized proposal based on the NN emulator is

$$\frac{1}{c_t} \hat{\pi}_t(\mathbf{x}) = \frac{1}{c_t} \sum_{i=1}^{J_t} v_i p_i(\mathbf{x}) = \sum_{i=1}^{J_t} \bar{v}_i p_i(\mathbf{x}),$$

where

$$\bar{v}_i = \frac{v_i}{c_t} = \frac{\pi(\mathbf{x}_i) |\mathcal{R}_i|}{\sum_{j=1}^{J_t} \pi(\mathbf{x}_j) |\mathcal{R}_j|}, \quad i = 1, \dots, N,$$

are also normalized. In order to sample $\frac{1}{c_t} \hat{\pi}_t(\mathbf{x})$, we would first (i) draw an index i^* from the set $\{1, \dots, J_t\}$ with probabilities $\bar{v}_i = \frac{1}{c_t} v_i$ ($i = 1, \dots, J_t$), and then (ii) sample from $p_{i^*}(\mathbf{x})$. In practice, we do not know the measures $|\mathcal{R}_i|$ and we are not able to draw samples uniformly in \mathcal{R}_i . Hence, we use SIR method to solve the problem drawing from an auxiliary pdf $\bar{q}_{\text{aux}}(\mathbf{x})$ (see Appendix A), as we have proposed in RADIS. Namely, we resample from the set $\{\mathbf{z}_{t,\ell}\}_{\ell=1}^L \sim \bar{q}_{\text{aux}}(\mathbf{x})$ with probabilities proportional to $\gamma_{t,\ell} = \frac{\hat{\pi}_t(\mathbf{z}_{t,\ell})}{\bar{q}_{\text{aux}}(\mathbf{z}_{t,\ell})}$. Below, we consider the special case that $\bar{q}_{\text{aux}}(\mathbf{x})$ is uniform.

Approximating \bar{v}_i 's Let choose an uniform auxiliary density $\bar{q}_{\text{aux}}(\mathbf{x})$, i.e., $\bar{q}_{\text{aux}}(\mathbf{x}) = \frac{1}{|\mathcal{X}|}$ for all $\mathbf{x} \in \mathcal{X}$. We draw $\{\mathbf{z}_{t,\ell}\}_{\ell=1}^L$ from the uniform $\bar{q}_{\text{aux}}(\mathbf{x})$. Then, the IS weight associated with the ℓ -th sample is

$$\gamma_{t,\ell} \propto \hat{\pi}_t(\mathbf{z}_{t,\ell}) = \pi(\mathbf{x}_{k_\ell}),$$

where

$$\mathbf{x}_{k_\ell} = \arg \min_{\mathbf{x}_k \in \mathcal{S}_t} \|\mathbf{x}_k - \mathbf{z}_\ell\|,$$

i.e., \mathbf{x}_{k_ℓ} represents the NN of \mathbf{z}_ℓ within the set of J_t nodes. Consider now the i -th node \mathbf{x}_i . All samples whose NN is \mathbf{x}_i have weight proportional to $\pi(\mathbf{x}_i)$. We denote those samples as the set

$$\mathcal{U}_i = \{\mathbf{z}_{t,\ell} : \mathbf{x}_i = \arg \min_{\mathbf{x}_k} \|\mathbf{x}_k - \mathbf{z}_{t,\ell}\|\}. \quad (\text{B.4})$$

The number of samples within \mathcal{U}_i can be written as $|\mathcal{U}_i| = \sum_{\ell=1}^L \mathbb{I}(\mathbf{x}_{k_\ell} = \mathbf{x}_i)$. The probability of resampling a $\mathbf{z}_{t,\ell}$ that comes from \mathcal{U}_i is proportional to $|\mathcal{U}_i| \pi(\mathbf{x}_i)$ (since there are $|\mathcal{U}_i|$ samples with weight $\pi(\mathbf{x}_i)$). As $L \rightarrow \infty$, by the law of large numbers, we have these probabilities converge to the true ones

$$\frac{|\mathcal{U}_i| \pi(\mathbf{x}_i)}{\sum_{k=1}^{J_t} |\mathcal{U}_k| \pi(\mathbf{x}_k)} \rightarrow \frac{|\mathcal{R}_i| \pi(\mathbf{x}_i)}{\sum_{k=1}^{J_t} |\mathcal{R}_k| \pi(\mathbf{x}_k)} = \bar{v}_i. \quad (\text{B.5})$$

Rejection sampling Note also that the samples within \mathcal{U}_i form a particle approximation of the uniform density over \mathcal{R}_i . Indeed, taking one sample at random from \mathcal{U}_i corresponds to applying rejection sampling on $p_i(\mathbf{x})$. In order to see this, consider the rejection sampling setting where $p_i(\mathbf{x})$ is the target probability and $\bar{q}_{\text{aux}}(\mathbf{x}) = \frac{1}{|\mathcal{X}|}$ is the proposal. Note that $\frac{p_i(\mathbf{x})}{\bar{q}_{\text{aux}}(\mathbf{x})} = \frac{|\mathcal{X}|}{|\mathcal{R}_i|}$ for all $\mathbf{x} \in \mathcal{R}_i$, and $\frac{p_i(\mathbf{x})}{\bar{q}_{\text{aux}}(\mathbf{x})} = 0$ for all $\mathbf{x} \notin \mathcal{R}_i$, so $\bar{q}_{\text{aux}}(\mathbf{x})$ is a valid proposal for rejection sampling with rejection constant $M = \frac{|\mathcal{X}|}{|\mathcal{R}_i|}$ [2, Chapter 3]. In rejection sampling, we draw $\mathbf{z} \sim \bar{q}_{\text{aux}}(\mathbf{x})$, $u \sim \mathcal{U}[0, 1]$ and accept \mathbf{z} if

$$u \frac{|\mathcal{X}|}{|\mathcal{R}_i|} \bar{q}_{\text{aux}}(\mathbf{z}) \leq p_i(\mathbf{z}). \quad (\text{B.6})$$

If the condition holds, \mathbf{z} is an independent sample from $p_i(\mathbf{x})$. Otherwise we reject \mathbf{z} , draw another candidate \mathbf{z} and so on. Note that, when $\mathbf{z} \in \mathcal{R}_i$, we have

$$u \frac{|\mathcal{X}|}{|\mathcal{R}_i|} \frac{1}{|\mathcal{X}|} \leq \frac{1}{|\mathcal{R}_i|} \iff u \leq 1, \quad (\text{B.7})$$

so we always accept all \mathbf{z} 's that are closest to node \mathbf{x}_i , becoming i.i.d. samples from $p_i(\mathbf{x})$. Conversely, when $\mathbf{z} \notin \mathcal{R}_i$, we have the condition $u \leq 0$ that never holds, so that we always reject them. Namely, the set \mathcal{U}_i contains i.i.d. samples from $p_i(\mathbf{x})$, that have been obtained by rejection sampling.

Summary With the particular choice $\bar{q}_{\text{aux}}(\mathbf{x}) = \frac{1}{|\mathcal{X}|}$ for all $\mathbf{x} \in \mathcal{X}$, the SIR approach in NN-AIS is equivalent to (i) estimating by Monte Carlo the mixture probabilities \bar{v}_i , and (ii) applying rejection sampling to sample uniformly within each Voronoi region \mathcal{R}_i .

Appendix C. NN-AIS in unbonded domains

In this section, we recall how to extend the applicability of the nearest neighbor (NN) construction (see Sect. 4.2) when the domain \mathcal{X} is unbounded and show how to adapt the support of NN approximation.

C.1. NN-AIS with a fixed support in an unbounded domain

Consider again the following mixture proposal,

$$\varphi_t(\mathbf{x}) = \alpha_t \bar{q}_{\text{par}}(\mathbf{x}) + (1 - \alpha_t) \frac{1}{c_t} \hat{\pi}_t(\mathbf{x}), \quad (\text{C.1})$$

where $\alpha_t \in [0, 1]$ for all t , and $\bar{q}_{\text{par}}(\mathbf{x})$ is parametric pdf that covers properly the tails of the posterior π . Namely, $\bar{q}_{\text{par}}(\mathbf{x})$ is defined in the unbounded domain \mathcal{X} of π , whereas $\hat{\pi}_t(\mathbf{x})$ is built considering a bounded support $\mathcal{D} \subset \mathcal{X}$, decided in advance by the user.

Hence, φ_t is a valid proposal when using NN-AIS with unbounded \mathcal{X} . In this simple scenario, \mathcal{D} is fixed and does not vary with the iteration t . However, the information provided by the samples from $\tilde{q}_{\text{par}}(\mathbf{x})$ can be used to expand the support of $\hat{\pi}_t$, i.e., such it has an adaptive support, as described below.

C.2. Adapting support in NN-AIS

Let \mathcal{X} be unbounded and $\hat{\pi}_t$ be the surrogate model built with NN. Let $\mathcal{D}_t \subset \mathcal{X}$ denote the compact subset of \mathcal{X} where $\hat{\pi}_t$ is defined, i.e., $\hat{\pi}_t$ is zero outside \mathcal{D}_t . Note that \mathcal{D}_t depends on t . The set of current nodes \mathcal{S}_t is used to define the boundaries of \mathcal{D}_t . One possible way is as follows: Take \mathcal{D}_t as the hyperrectangle whose edges are defined by the maximum and minimum value, in each dimension, of the set \mathcal{S}_t , i.e.,

$$\mathcal{D}_t = \{\mathbf{x} \in \mathcal{X} : \min_{\mathbf{s}_{t-1} \in \mathcal{S}_{t-1}} s_{d,t-1} \leq x_d \leq \max_{\mathbf{s}_{t-1} \in \mathcal{S}_{t-1}} s_{d,t-1}, \\ d = 1, \dots, d_x\}, \quad (\text{C.2})$$

where x_d denotes the d -th element of \mathbf{x} , $\mathbf{s}_{t-1} = [s_{1,t-1}, \dots, s_{d_x,t-1}] \in \mathcal{S}_{t-1}$ and \mathcal{S}_{t-1} denotes the set of nodes at iteration t . After adding new nodes, we update the bounds of the hyperrectangle. Note that only samples from $\tilde{q}_{\text{par}}(\mathbf{x})$ that fall outside \mathcal{D}_t will expand it. Note that, the size of \mathcal{D}_t is always increasing but controlled by the tail of $\tilde{\pi}$. Indeed, the candidate samples drawn in the tails of $\tilde{\pi}$ will have very low values of π , so that the probability of sampling those regions will be negligible (then these regions will be never used). In the case we use a uniform $\tilde{q}_{\text{aux}}(\mathbf{x})$ in \mathcal{D}_t to sample $\hat{\pi}_t$, note that $\tilde{q}_{\text{aux}}(\mathbf{x})$ actually depends on t and is changing at every iteration whenever \mathcal{D}_t changes.

References

- [1] A. O'Hagan, Bayesian inference, in: Kendall's Advanced Theory of Statistics, vol. 2B, Arnold, London, UK, 1994.
- [2] L. Martino, D. Luengo, J. Miguez, Independent Random Sampling Methods, Springer, 2018.
- [3] C.P. Robert, G. Casella, Monte Carlo Statistical Methods, Springer, 2004.
- [4] M.F. Bugallo, V. Elvira, L. Martino, D. Luengo, J. Miguez, P.M. Djuric, Adaptive importance sampling: the past, the present, and the future, IEEE Signal Process. Mag. 34 (4) (2017) 60–79.
- [5] O. Cappé, A. Guillin, J.M. Marin, C.P. Robert, Population Monte Carlo, J. Comput. Graph. Stat. 13 (4) (2004) 907–929.
- [6] O. Cappé, R. Douc, A. Guillin, J.M. Marin, C.P. Robert, Adaptive importance sampling in general mixture classes, Stat. Comput. 18 (2008) 447–459.
- [7] V. Elvira, L. Martino, D. Luengo, M.F. Bugallo, Improving population Monte Carlo: alternative weighting and resampling schemes, Signal Process. 131 (2017) 77–91.
- [8] Y. El-Laham, P.M. Djuric, M.F. Bugallo, A variational adaptive population importance sampler, in: ICASSP 2019–2019 IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP, IEEE, 2019, pp. 5052–5056.
- [9] L. Martino, V. Elvira, D. Luengo, J. Corander, An adaptive population importance sampler: learning from the uncertainty, IEEE Trans. Signal Process. 63 (16) (2015) 4422–4437.
- [10] A. O'Hagan, Bayesian analysis of computer code outputs: a tutorial, Reliab. Eng. Syst. Saf. 91 (10–11) (2006) 1290–1300.
- [11] D. Busby, Hierarchical adaptive experimental design for Gaussian process emulators, Reliab. Eng. Syst. Saf. 94 (7) (2009) 1183–1193.
- [12] C.J. Schmit, J.R. Pritchard, Emulation of reionization simulations for Bayesian inference of astrophysics parameters using neural networks, Mon. Not. R. Astron. Soc. 475 (1) (2018) 1213–1223.
- [13] S. Razavi, B.A. Tolson, D.H. Burn, Review of surrogate modeling in water resources, Water Resour. Res. 48 (7) (2012).
- [14] D.H. Svendsen, L. Martino, G. Camps-Valls, Active emulation of computer codes with Gaussian processes - application to remote sensing, Pattern Recognit. 100 (2020) 107103.
- [15] J.A. Christen, C. Fox, Markov Chain Monte Carlo using an approximation, J. Comput. Graph. Stat. 14 (4) (2005) 795–810.
- [16] H. Ying, K. Mao, K. Mosegaard, Moving target Monte Carlo, arXiv preprint, arXiv:2003.04873.
- [17] P.R. Conrad, Y.M. Marzouk, N.S. Pillai, A. Smith, Accelerating asymptotically exact MCMC for computationally intensive models via local approximations, J. Am. Stat. Assoc. 111 (516) (2016) 1591–1607.
- [18] J.S. Liu, Monte Carlo Strategies in Scientific Computing, Springer Science & Business Media, 2008.
- [19] M. Kennedy, Bayesian quadrature with non-normal approximating functions, Stat. Comput. 8 (4) (1998) 365–375.
- [20] C.E. Rasmussen, Z. Ghahramani, Bayesian Monte Carlo, in: Advances in Neural Information Processing Systems, 2003, pp. 505–512.
- [21] F. Llorente, L. Martino, V. Elvira, D. Delgado, J. Lopez-Santiago, Adaptive quadrature schemes for Bayesian inference via active learning, IEEE Access 8 (2020) 208462–208483.
- [22] E. Cleary, A. Garbuno-Inigo, S. Lan, T. Schneider, A.M. Stuart, Calibrate, emulate, sample, arXiv:2001.03689.
- [23] W.R. Gilks, N.G. Best, K.K.C. Tan, Adaptive rejection Metropolis sampling within Gibbs sampling, Appl. Stat. 44 (4) (1995) 455–472.
- [24] R. Meyer, B. Cai, F. Perron, Adaptive rejection Metropolis sampling using Lagrange interpolation polynomials of degree 2, Comput. Stat. Data Anal. 52 (7) (2008) 3408–3423.
- [25] L. Martino, J. Read, D. Luengo, Independent doubly adaptive rejection Metropolis sampling within Gibbs sampling, IEEE Trans. Signal Process. 63 (12) (2015) 3123–3138.
- [26] L. Martino, H. Yang, D. Luengo, J. Kannianen, J. Corander, A fast universal self-tuned sampler within Gibbs sampling, Digit. Signal Process. 47 (2015) 68–83.
- [27] L. Martino, R. Casarin, F. Leisen, D. Luengo, Adaptive independent sticky MCMC algorithms, EURASIP J. Adv. Signal Process. 2018 (1) (2018) 5.
- [28] W.R. Gilks, P. Wild, Adaptive rejection sampling for Gibbs sampling, Appl. Stat. 41 (2) (1992) 337–348.
- [29] W. Hörmann, A rejection technique for sampling from T-concave distributions, ACM Trans. Math. Softw. 21 (2) (1995) 182–193.
- [30] D. Görür, Y.W. Teh, Concave convex adaptive rejection sampling, J. Comput. Graph. Stat. 20 (3) (2011) 670–691.
- [31] G. Marsaglia, W.W. Tsang, The Ziggurat method for generating random variables, J. Stat. Softw. 8 (5) (2000) 1–7.
- [32] J. Felip, N. Ahuja, O. Tickoo, Tree pyramidal adaptive importance sampling, arXiv preprint, arXiv:1912.08434.
- [33] G.P. Lepage, A new algorithm for adaptive multidimensional integration, J. Comput. Phys. 27 (2) (1978) 192–203.
- [34] J.H. Friedman, M.H. Wright, A nested partitioning procedure for numerical multiple integration, ACM Trans. Math. Softw. 7 (1) (1981) 76–92.
- [35] W.H. Press, G.R. Farrar, Recursive stratified sampling for multidimensional Monte Carlo integration, Comput. Phys. 4 (2) (1990) 190–195.
- [36] X. Lu, T. Rainforth, Y. Zhou, J.-W. van de Meent, Y.W. Teh, On exploration, exploitation and learning in adaptive importance sampling, arXiv preprint, arXiv:1810.13296.
- [37] T. Foster, C.L. Lei, M. Robinson, D. Gavaghan, B. Lambert, Model evidence with fast tree based quadrature, arXiv preprint, arXiv:2005.11300.
- [38] T.E. Hanson, J.V.D. Monteiro, A. Jara, The Polya tree sampler: toward efficient and automatic independent Metropolis–Hastings proposals, J. Comput. Graph. Stat. 20 (1) (2011) 41–62.
- [39] P. Zhang, Nonparametric importance sampling, J. Am. Stat. Assoc. 91 (435) (1996) 1245–1253.
- [40] R.J. Steele, A.E. Raftery, M.J. Emond, Computing normalizing constants for finite mixture models via Incremental Mixture Importance Sampling (IMIS), J. Comput. Graph. Stat. 15 (3) (2006) 712–734.
- [41] L. Martino, V. Elvira, D. Luengo, J. Corander, Layered adaptive importance sampling, Stat. Comput. 27 (3) (2017) 599–623.
- [42] V. Elvira, L. Martino, D. Luengo, M.F. Bugallo, Generalized multiple importance sampling, Stat. Sci. 34 (1) (2019) 129–155.
- [43] J.M. Cornuet, J.M. Marin, A. Mira, C.P. Robert, Adaptive multiple importance sampling, Scand. J. Stat. 39 (4) (2012) 798–812.
- [44] M. Balesdent, J. Morio, J. Marzat, Kriging-based adaptive importance sampling algorithms for rare event estimation, Struct. Saf. 44 (2013) 1–10.
- [45] V. Dubourg, B. Sudret, F. Deheeger, Metamodel-based importance sampling for structural reliability analysis, Probab. Eng. Mech. 33 (2013) 47–57.
- [46] R.M. Neal, Annealed importance sampling, Stat. Comput. 11 (2) (2001) 125–139.
- [47] V. Elvira, L. Martino, P. Ciosas, Importance Gaussian quadrature, arXiv:2001.03090.
- [48] F. Llorente, L. Martino, D. Delgado, J. Lopez-Santiago, Marginal likelihood computation for model selection and hypothesis testing: an extensive review, arXiv:2001.0052.
- [49] L. Martino, J. Read, V. Elvira, F. Louzada, Cooperative parallel particle filters for on-line model selection and applications to urban mobility, Digit. Signal Process. 60 (2017) 172–185.
- [50] H. Niederreiter, Random Number Generation and Quasi-Monte Carlo Methods, Society for Industrial Mathematics, 1992.
- [51] J.S. Liu, Monte Carlo Strategies in Scientific Computing, Springer, 2004.
- [52] L. Martino, V. Elvira, F. Louzada, Effective sample size for importance sampling based on discrepancy measures, Signal Process. 131 (2017) 386–401.
- [53] Ö.D. Akyildiz, J. Miguez, Convergence rates for optimised adaptive importance samplers, arXiv preprint, arXiv:1903.12044.

- [54] D.B. Rubin, Using the SIR algorithm to simulate posterior distributions, in: Bernardo, Degroot, Lindley, Smith (Eds.), *Bayesian Statistics 3*, Oxford University Press, Oxford, 1988.
- [55] A.F. Smith, A.E. Gelfand, Bayesian statistics without tears: a sampling-resampling perspective, *Am. Stat.* 46 (2) (1992) 84–88.
- [56] G.H. Givens, J.A. Hoeting, *Computational Statistics*, vol. 703, John Wiley & Sons, 2012.
- [57] C.P. Robert, W. Changye, Markov Chain Monte Carlo methods, a survey with some frequent misunderstandings, arXiv preprint, arXiv:2001.06249.
- [58] L. Martino, V. Elvira, G. Camps-Valls, Group importance sampling for particle filtering and MCMC, *Digit. Signal Process.* 82 (2018) 133–151.
- [59] C. Musso, N. Oudjane, F. Le Gland, Improving regularised particle filters, in: A. Doucet, N. de Freitas, N. Gordon (Eds.), *Sequential Monte Carlo Methods in Practice*. Statistics for Engineering and Information Science, Springer, New York, 2001, pp. 247–271.
- [60] E. Veach, L.J. Guibas, Optimally combining sampling techniques for Monte Carlo rendering, in: *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*, 1995, pp. 419–428.
- [61] C.E. Rasmussen, C.K.I. Williams, *Gaussian Processes for Machine Learning*, The MIT Press, New York, 2006.
- [62] L. Martino, J. Read, Joint introduction to Gaussian processes and relevance vector machines with connections to Kalman filtering and other kernel smoothers, arXiv:2009.09217.
- [63] M. Osborne, R. Garnett, Z. Ghahramani, D.K. Duvenaud, S.J. Roberts, C.E. Rasmussen, Active learning of model evidence using Bayesian quadrature, in: *Advances in Neural Information Processing Systems*, 2012, pp. 46–54.
- [64] L. Martino, Parsimonious adaptive rejection sampling, *Electron. Lett.* 53 (16) (2017) 1115–1117.
- [65] A. O'Hagan, Probabilistic uncertainty specification: overview, elaboration techniques and their application to a mechanistic model of carbon flux, *Environ. Model. Softw.* 36 (2012) 35–48, <https://doi.org/10.1016/j.envsoft.2011.03.003>, thematic issue on Expert Opinion in Environmental Modelling and Management.
- [66] M. Kennedy, A. O'Hagan, Bayesian calibration of computer models, *J. R. Stat. Soc., Ser. B, Stat. Methodol.* 63 (3) (2001) 425–450.
- [67] G. Camps-Valls, D. Sejdinovic, J. Runge, M. Reichstein, A perspective on Gaussian processes for Earth observation, *Nat. Sci. Rev.* 6 (2019) 616–618.
- [68] X. Dong, D. Thanou, M. Rabbat, P. Frossard, Learning graphs from data: a signal representation perspective, *IEEE Signal Process. Mag.* 36 (3) (2019) 44–63.
- [69] A. Owen, Y. Zhou, Safe and effective importance sampling, *J. Am. Stat. Assoc.* 95 (449) (2000) 135–143.
- [70] P.C. Gregory, Bayesian re-analysis of the Gliese 581 exoplanet system, *Mon. Not. R. Astron. Soc.* 415 (3) (2011) 2523–2545, <https://doi.org/10.1111/j.1365-2966.2011.18877.x>, arXiv:1101.0800.
- [71] S.C.C. Barros, et al., WASP-113b and WASP-114b, two inflated hot Jupiters with contrasting densities, *Astron. Astrophys.* 593 (2016) A113, <https://doi.org/10.1051/0004-6361/201526517>.
- [72] T. Trifonov, S. Stock, T. Henning, S. Reffert, M. Kürster, M.H. Lee, B. Bitsch, R.P. Butler, S.S. Vogt, Two Jovian planets around the Giant Star HD 202696: a growing population of packed massive planetary pairs around massive stars?, *Astron. J.* 157 (3) (2019) 93, <https://doi.org/10.3847/1538-3881/aafa11>, arXiv:1901.01935.
- [73] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, *Numerical Recipes in C++: The Art of Scientific Computing*, Springer, 2002.
- [74] S. Jacquemoud, W. Verhoef, F. Baret, C. Bacour, P. Zarco-Tejada, G. Asner, C. François, S. Ustin, PROSPECT+ SAIL models: a review of use for vegetation characterization, *Remote Sens. Environ.* 113 (2009) S56–S66.
- [75] G. Camps-Valls, D. Svendsen, L. Martino, J. Munoz-Mari, V. Laparra, M. Campos-Taberner, D. Luengo, Physics-aware Gaussian processes in remote sensing, *Appl. Soft Comput.* 68 (2018) 69–82, <https://doi.org/10.1016/j.asoc.2018.03.021>.
- [76] A. Gelman, X.-L. Meng, *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, John Wiley & Sons, 2004.
- [77] T. Butler, L. Graham, S. Mattis, S. Walsh, A measure-theoretic interpretation of sample based numerical integration with applications to inverse and prediction problems under uncertainty, *SIAM J. Sci. Comput.* 39 (5) (2017) A2072–A2098.
- [78] L. Devroye, L. Györfi, G. Lugosi, H. Walk, On the measure of Voronoi cells, *J. Appl. Probab.* 54 (2) (2017) 394–408.
- [79] A. Stuart, A. Teckenstrup, Posterior consistency for Gaussian process approximations of Bayesian posterior distributions, *Math. Comput.* 87 (310) (2018) 721–753.



Fernando Llorente received a B.Sc. in Physics from Universidad Autónoma de Madrid in 2016 and a M.Sc. in Mathematical Engineering from Universidad Carlos III de Madrid in 2018. He is currently doing his PhD in the Department of Statistics at the Universidad Carlos III de Madrid, where his main research interests are Bayesian inference and statistical signal processing.



Luca Martino was born in Palermo, Italy, 1980. He received his M.Sc. degree in electronic engineering from the Politecnico di Milano, Italy, and obtained his PhD in Statistical Signal Processing from the Universidad Carlos III de Madrid, Spain, in 2011. Later, he spent two years in the Department of Statistics at the University of Helsinki, Finland. He carried out a postdoc research at Sao Paulo Research Foundation FAPESP, and at Universitat de Valencia. He was assis-

tant professor at Universidad Carlos III de Madrid from September 2017 to December 2018. Currently, he is Associate Professor at Universidad Rey Juan Carlos.



David Delgado received the M.Sc degree in mathematics from the Autonomous University of Madrid, Madrid, Spain, in 1999 and the Ph.D. degree in applied mathematics from the Technical University of Denmark, Lyngby, in 2005. He combines research and teaching as Associate Professor at Universidad Carlos III de Madrid. His research interests include applied multivariate statistical analysis, machine learning, and pattern recognition. He has published 39 JCR articles

and he has received several teaching and research awards.



Gustau Camps Valls received a PhD in Physics in 2002 from the Universitat de València and he is currently Full professor in Electrical Engineering, and coordinator in the Image and Signal Processing (ISP) group, <http://isp.uv.es>, at the Universitat de València. He is interested in the development of machine learning algorithms for geoscience and remote sensing data analysis.